PESI 7/24/07 2:15 PM

WAGE RISK AND EMPLOYMENT RISK OVER THE LIFE CYCLE

Hamish Low
Costas Meghir
Luigi Pistaferri

THE INSTITUTE FOR FISCAL STUDIES
WP06/27

Wage Risk and Employment Risk over the Life Cycle*

Hamish Low[†], Costas Meghir[‡], and Luigi Pistaferri.[§]

January 2007

Abstract

This paper decomposes the sources of risk to income that individuals face over their lifetimes. We distinguish productivity risk from employment risk and identify the components of each using the Survey of Income and Program Participation and the Panel Study of Income Dynamics. Estimates of productivity risk controlling for employment risk and for individual labour supply choices are substantially lower than estimates that attribute all wage variation to productivity risk. We use a partial equilibrium life-cycle model of consumption and labour supply to analyse the choices individuals make in the light of these risks and to measure the welfare cost of the different types of risk. Productivity risk induces a considerably greater welfare loss than employment risk primarily because productivity shocks are more persistent. Reflecting this, the welfare value of government programs such as food stamps which partially insure productivity risk is greater than the value of unemployment insurance which provides (partial) insurance against employment risk and no insurance against persistent shocks.

Keywords: uncertainty, life-cycle models, unemployment, precautionary savings

JEL Classification: D91, H31, J64

*We are grateful to Joe Altonji, Gadi Barlevy, Marco Bassetto, Audra Bowlus, Thomas Crossley, James Heckman, Bruce Meyer, Dale Mortensen, Giuseppe Moscarini, Chris Pissarides, Victor Rios-Rull, Chris Taber, Robert Townsend, Gianluca Violante, conference participants at the NBER Summer Institute, Society of Economic Dynamics, Stanford Institute for Theoretical Economics, the 2006 Cowles Macro-Labor Conference, the "Macroeconomics of Imperfect Risk Sharing" conference in Santa Barbara, as well as participants in numerous seminars and workshops. Pistaferri acknowledges funding from NSF through grant 0453119. We thank the ESRC Centre for Microeconomic Analysis of Public Policy at the Institute for Fiscal Studies.

[†]Cambridge University and IFS.

[‡]University College London and IFS.

[§]Stanford University, SIEPR and CEPR.

Executive Summary

This paper decomposes the source of the risk to income that individuals face over their lifetimes. We distinguish between unemployment risk and wage risk. Unmployment risk encompasses a number of components: the uncertainty individuals face about receiving a job offer while unemployed; the uncertainty about whether a particular job will continue or a firm close; the uncertainty about the type of firm that an individual is matched with. Wage risk exists independently of the employer's characteristics, rather it is the uncertainty an individual faces over their own productivity. Shocks to productivity persist across matches with different firms, in contrast to unemployment which do not persist across different jobs.

The different sources of risk matter for two main reasons: The first reason is that risks differ in their insurance opportunities. For example, layoff risk is often partially insured by the unemployment insurance system, while individual wage risk is rarely insured in any formal way. It is precisely this lack of formal insurance that prompts prudent individuals to engage in precautionary behavior. Furthermore, the individual's response to wage risk will depend partly on the availability of outside insurance - private or public. The second reason is that it is important to distinguish between movements in earnings that reflect choice and those which reflect uncertainty. We address this issue by allowing for a choice over labor supply and allowing for job mobility which implies that a proportion of earnings fluctuations, usually interpreted as risk, are in fact attributed to choice.

The contribution of this paper is to provide a life-cycle framework for making a meaningful distinction between the different sorts of risk that people face (productivity risk and employment risk), and then to estimate the extent of risk within this framework. This enables us to show how individuals respond to the different types of risk in a calibrated life-cycle model, to calculate the welfare costs of risk, and to evaluate the effect of various government insurance programs.

We find the following results. If mobility is ignored, all the wage variability that is due to matching effects is attributed to permanent shocks. Since job mobility is valued because of the absence of a downside (all bad offers can be rejected), this biases upwards the amount of permanent uncertainty. Welfare calculations of the risk premium show that both high and low educated individuals are willing to pay considerably more to avoid productivity risk than employment risk. We then evaluate the welfare value of three insurance programs we have included in the model: Unemployment Insurance, Disability Insurance and a universal means-tested program. We find that, of these, the latter is by far the most valuable. This is because of the seriousness of productivity risk which derives from the permanence of productivity shocks.

1 Introduction

There is now an extensive literature analyzing individuals' precautionary response to income risk under incomplete markets. The theoretical literature has made enormous progress in clarifying under which circumstances precautionary behavior arises, and in throwing some light on the role of income uncertainty, income persistence, and the degree of market incompleteness.[1] Empirical analysis has concentrated on assessing the levels of income risk and measuring its effects on life-cycle consumption profiles and wealth accumulation.[2] However, in most earlier studies labor supply is taken as exogenous and there is a unique source of risk attributed to income fluctuations.[3] Here we generalize the model in both directions. First we allow labor supply to be endogenous with individuals deciding whether to work or not depending on the returns to work, thus introducing an additional mechanism for self insurance. Second we allow individuals to make endogenous job mobility decisions in a world where wages depend on match specific effects. This now allows us to introduce a distinction between employment risk and productivity risk.[4] Within this framework we are able to discuss the welfare costs of the different sources of (uninsurable) risk and to quantify the value that individuals attribute to key welfare programs.

The different sources of risk matter for two main reasons: First, risks differ in their insurance opportunities. For example, layoff risk is often partially insured by the unemployment insurance system, while individual productivity risk is rarely insured in any formal way because of moral hazard and limited enforcement and commitment reasons. It is precisely this lack of formal insurance that prompts prudent individuals to engage in precautionary behavior. Furthermore, the individual's response to earnings risk will depend partly on the availability of outside insurance - private or public. With few exceptions (Hubbard, Skinner and Zeldes, 1995), the literature on precautionary savings has ignored this and assumed that only self-insurance is available. In this paper, we propose a model

[1] See Kimball (1990), Caballero (1990), and Deaton (1991).

[2] On the measurement of the level of risk, see MaCurdy (1982), Abowd and Card (1989), and Meghir and Pistaferri (2004). For the implications of the effects of this risk, see amongst others Carroll and Samwick (1997), Banks, Blundell and Brugiavini (2001), Zeldes (1986), Deaton (1991), Carroll (1992), Gourinchas and Parker (2002), and Attanasio, Banks, Meghir and Weber (1995).

[3] Exceptions are French (2005) and Low (2005).

[4] Low (2005) analyzes the joint saving and labour supply decision, but in a context without exogenous job destruction or search frictions. Lentz (2003) analyzes the interaction between search frictions and saving, but ignores the risk to individuals' own productivity which is independent of any particular match. See also Costain (1999) for an equilibrium search model with precautionary savings that attempts to measure the welfare effects of unemployment insurance. Finally, Rendon (2006) examines the relationship between wealth accumulation and job search dynamics in a model where the motivation for accumulating wealth is to finance voluntary quits in order to search for a better job.

in which people can self-insure, but may also be eligible for government provided insurance mirroring three popular programs in the US: Unemployment Insurance (UI), Disability Insurance (DI), and Food Stamps. At retirement, people collect social security benefits (pensions). Unemployment insurance is aimed at insuring against exogenous job destruction and (partly) against the difficulty of finding a new job. The disability insurance system is supposed to provide insurance against an extreme form of productivity shock which results in permanent inability to work. Finally, the Food Stamps program provides *universal* insurance against low income, whatever its cause. It is worth stressing that all these systems provide partial insurance only.

Second, it is important to distinguish between movements in earnings that reflect choice and those which reflect uncertainty. We address this issue by allowing for endogenous labor supply and job mobility which implies that a proportion of earnings fluctuations, usually interpreted as risk, are in fact attributed to choice.[5] Having removed the effect of these sources of fluctuations, we also decompose earnings fluctuations into transitory shocks (which we assume are entirely attributable to measurement error) and permanent shocks (which we take as a source of uncertainty).

Thus the contribution of this paper is to provide a life-cycle framework for making a meaningful distinction between the different sorts of risk that people face (productivity risk and employment risk), and then to estimate the extent of risk within this framework using longitudinal wage and job mobility data from the Survey of Income and Program Participation (SIPP) and unemployment duration data from the Panel Study of Income Dynamics (PSID). This enables us to show how individuals respond to the different types of risk in a calibrated life-cycle model of intertemporal consumption and labor supply, to calculate the welfare costs of risk allowing for the various substitution effects, and to evaluate the effect of various government insurance programs allowing for some of the moral hazard distortions they induce.

Consider first the distinction between productivity risk and employment risk. In our model productivity risk is individual-specific uncertainty which exists independently of the employer's characteristics. We follow the empirical evidence on wage dynamics and assume that productivity shocks result in permanent shifts of the wage profile. Unemployment risk captures the uncertainty about

[5] More generally, we follow the standard approach of using "realizations" to infer risk, thereby assuming that the individual's and the econometrician's information set coincide. Primarily for lack of adequate data, we abstract from the important issues that have to do with consumers having superior information vis-á-vis the econometrician (see Manski, 2004, for a discussion, and Pistaferri, 2001, 2003, for empirical examples).

having a job and also about the firm type. This includes the possibility of firm closure or job destruction, the difficulty of finding a new job match while unemployed, and the extent of unobserved heterogeneity across firms. In a fully competitive labor market with no worker-firm match heterogeneity and no search costs the distinction between employment and wage/productivity risk is meaningless: In that world individual wages may fluctuate due to unexpected shocks, but given wages individuals decide whether they wish to work or not; unemployment is not a source of risk. It is the interaction of shocks to individuals or firms (job destruction) with search costs and firm heterogeneity that give rise to the distinction between these risks.

To implement our model we estimate a number of parameters capturing the risks that people face and then simulate behavior as the magnitude of the various risks change. The parameters of interest for our simulations are obtained partly from estimating the characteristics of the wage dynamic process with endogenous participation and mobility choices, and partly from calibrating our life-cycle model to fit observed participation profiles and unemployment durations. All parameters reflect important heterogeneity by skills (here captured by schooling).

In addressing the question of how individuals respond to risk, we begin by simulating savings and participation behavior for individuals facing the estimated risks. These simulations give an indication of the extent of precautionary behavior (both precautionary saving and precautionary labor supply). We then calculate how much individuals would be willing to pay to avoid the various risks, how much of the precautionary response is due to employment risk and how much to wage/productivity risk. Finally, we measure the value that people attach to the various government provided insurance programs in our model. There is a clear relationship between the results on the costs of each source of risk and the value of the government programs since these programs are designed to insure different types of risk.

We find the following results. If mobility is ignored, all the wage variability that is due to matching effects is attributed to permanent shocks. Since job mobility is valued because of the absence of a downside (all bad offers can be rejected), this biases upwards the amount of permanent uncertainty by a factor of 40% and leads to upward bias in the amount of precautionary saving people hold against permanent productivity risk.[6] Welfare calculations of the risk premium show that both high

[6] There would be no bias if job mobility were completely at random.

and low educated individuals are willing to pay considerably more to avoid productivity risk than employment risk. As a result, in assessing the reasons for holding assets for precautionary reasons, productivity risk dominates. However, ignoring employment risk leads to inaccurate predictions about job and unemployment durations.

We then evaluate the welfare value of three insurance programs we have included in the model: Unemployment Insurance, Disability Insurance and a universal means-tested program. We find that, of these, the latter is by far the most valuable. This is because of the importance of productivity risk which derives from the permanence of productivity shocks.

The layout of the paper is as follows. Section 2 presents the model and discusses the distinction between employment and productivity risk. Section 3 describes the data. Section 4 describes the estimation strategy and results for estimating the wage process. Section 5 presents the calibration process for the remaining parameters. Section 6 presents our calculations of the welfare costs of uncertainty and the welfare benefit of government insurance programs. Section 7 concludes.

2 Model

2.1 Overview

We specify a model where individuals choose consumption and make work decisions so as to maximize an intertemporal utility function, in an environment with search frictions. Individuals face multiple sources of uncertainty: in each period employed individuals may be laid off or may receive offers of alternative employment; unemployed workers may or may not be offered a job; all individuals face uninsurable shocks to their productivity. The economy offers partial social insurance in the form of a number of programs. These are Food Stamps, Unemployment Insurance, Disability Insurance and Social Security (pensions). Any change to these programs is funded by proportional taxation; thus individuals are linked through the government budget constraint. The model has numerous sources of dynamics. These include asset accumulation, the fact that job offer probabilities are state dependent and that current actions affect future eligibility to the various programs. We consider two types of individuals separately: the lower educated individuals which include all those with a high school diploma or less and the higher educated individuals with at least some College.

2.2 Structure of Wages and Shocks

We begin the model specification by outlining the process for wages. We assume that wages w_{it} in the data are governed by the process:

$$\ln w_{it} = d_t^{ed} + x_{it}'\psi^{ed} + u_{it} + e_{it} + a_{ij(t_0)} \tag{1}$$

where w_{it} is the real hourly wage, d_t^{ed} represents the log price of human capital at time t for education group ed, x_{it} a vector of regressors including age, ψ^{ed} is a vector of parameters specific to the individual's education group, u_{it} the permanent component of wages, and e_{it} measurement error.

The term $a_{ij(t_0)}$ denotes a firm-worker match specific component where $j(t_0)$ indexes the firm that the worker joined in period $t_0 \leq t$.[7] It is drawn from a normal distribution with mean zero and variance σ_a^2. We model the match effect as constant over the life of the worker-employer relationship, and so if the worker does not change employer between t and $t+1$, there is no wage growth due to the match effect. If the worker switches to a different employer between t and $t+1$, however, there will be some wage growth which we can term a mobility premium. In this case we define the random variable $\xi_{it+1} = a_{ij(t+1)} - a_{ij(t_0)}$ as the wage growth due to inter-firm mobility between t and $t+1$. Since offers can be rejected when received, only a censored distribution of ξ_{it+1} is observed. The match effect $a_{ij(.)}$ is complementary to individual productivity. It is constant over time but it will be assumed uncertain across firms.[8] Both the match effect and the idiosyncratic shock have education specific distributions. The information structure is such that workers and firms are completely informed about u_{it} and $a_{ij(.)}$ when they meet (jobs are "search goods"). The importance of match effects in explaining wages has been stressed by Topel and Ward (1992) and Abowd, Kramarz and Margolis (1999). Postel-Vinay and Robin (2002) show in an equilibrium setting how firm and individual heterogeneity translate into a match effect. In our setting we do not allow the firm to respond to outside offers. Finally, we assume that there are constant returns to scale in labor implying that the firm is willing to hire anyone who can produce non-negative rents.

[7] We should formally have a j subscript on wages but since it does not add clarity we have dropped it. Note also that in the absence of firm data one cannot distinguish between a pure firm effect and a pure match effect. In the latter case, one can imagine $\alpha_{ij(t_0)}$ as being the part of the matching rent that accrues to the worker. We take the bargaining process that produces this sharing outcome as given.

[8] Ideally we would like to allow for shocks to the match effect. These will act as within firm aggregate shocks and their nature. Restricting match effects to be constant is forced upon us by the lack of matched firm and individual data. It is however, possible to extend the model in this interesting way when suitable data becomes available.

Following a number of papers in the literature[9] we assume that the permanent component of wages follows a random walk process:[10]

$$u_{it} = u_{it-1} + \zeta_{it} \tag{2}$$

where ζ_{it} is a normally distributed random shock with mean zero and variance σ_ζ^2. We take this shock to be uncertain and variable from period to period.[11] Given a particular level of unobserved productivity, the worker will be willing to work for some firms but not for others, depending on the value of the match.

Identifying the variance of a transitory shock from that of measurement error is generally not possible without further assumptions or without results from a validation study. Here we assume that e_{it} consists entirely of measurement error and is normally distributed with variance σ_e^2.

As far as the policy implications of the model are concerned we are interested in estimating σ_a^2 and σ_ζ^2. We describe later how these are identified and estimated.

2.3 Job destruction and job arrival rates

In each period each worker will receive an alternative job offer with probability λ^e. Those who are currently unemployed receive an offer with probability λ^n. Individuals become unemployed either because they choose to quit following particular wage realizations or because of exogenous job destruction, which happens each period with probability δ. The friction parameters as well as the variance parameters are all assumed to be specific to an education group.

We assume there is no exogenous depreciation of skills following job loss. Instead, the loss of the particular match on entering unemployment may lead to wages on re-entry being lower because the new firm will on average have a lower match value. This is the case because individuals in work will

[9] See MaCurdy (1982), Topel (1991), Abowd and Card (1989), Moffitt and Gottschalk (2001), and Meghir and Pistaferri (2004) .

[10] It is possible that observed wages may have already been smoothed out relative to productivity by implicit agreements within the firm. This means that productivity risk may be greater than observed wage movements within a firm, which implies that the process for productivity shocks is not properly identified for the unemployed. In other words, productivity shocks are a combination of actual shocks plus insurance, but this insurance is only present if the individual is working. If the unemployed experience greater productivity risk than estimated, this will impact on the reservation wage and on job search. For the time being we ignore this issue as far as permanent shocks are concerned. On the other hand we ignore transitory shocks to wages (the component e_{it} in (1) is assumed to reflect measurement error).

[11] Farber and Gibbons (1996) assume that individual productivity is unknown to the firm, but it is learned over time through observation of output, and so wages are updated in a Bayesian sense. They prove that this will result in the wage residual being a martingale. Thus our unit root characterization can also be consistent with a less than complete information case, but we have not considered the implications of the learning case as yet.

have improved over the average offer through job mobility, before their job is destroyed.[12] This is especially likely for exogenous job destruction. Thus firm heterogeneity implies that exogenous job destruction will lead to apparent scarring.

2.4 Individual Optimization

We consider an individual with a period utility function

$$U_t = U(c_t, P_t)$$

where P_t is a discrete $\{0, 1\}$ labor supply participation variable and c_t consumption. The individual is assumed to maximize lifetime expected utility,

$$\max_{c,P} V_{it} = E_t \sum_{s=t}^{L} \beta^{s-t} U(c_{is}, P_{is})$$

where β is the discount factor and E_t the expectations operator conditional on information available in period t (a period being a quarter of a year). Individuals live for L periods, may work from age 22 to 62, and face an exogenous mandatory spell of retirement of 10 years at the end of life. The date of death is known with certainty.

The worker's problem is to decide whether to work or not and, if the opportunity arises, whether to switch firm. When unemployed he has to decide whether to accept a job that may have been offered or wait longer. If eligible, the unemployed person will have the option to apply for disability insurance. There is a fixed, known probability of being successful, conditional on applying. Whether employed or not, the individual has to decide how much to save and consume. Accumulated savings can be used to finance spells out of work and early retirement.

We use a utility function of the form

$$U(c, P) = \frac{(c \cdot \exp\{\eta P\})^{1-\gamma}}{1-\gamma}$$

We consider cases where $\gamma > 1$ and $\eta < 0$, implying that individuals are reasonably risk averse, participation reduces utility and that consumption and participation are Frisch complements (i.e. the marginal utility of consumption is higher when participating).

[12] Indeed, as stated by Jacubson, LaLonde and Sullivan (1993), "workers possessing skills that were especially suited to their old positions are likely to be less productive, at least initially, in their subsequent jobs. Such a fit between workers' skills and the requirements of their old jobs could have resulted from on-the-job investment in firm-specific human capital or from costly search resulting in particularly good match with their old firms".

The intertemporal budget constraint during the working life has the form

$$A_{it+1} = R\left[A_{it} + \left(w_{it}h\left(1-\tau_w\right) - F_{it}\right)P_{it} + \left(B_{it}E_{it}^{UI}\left(1-E_{it}^{DI}\right) + D_{it}E_{it}^{DI}\right)\left(1-P_{it}\right) + T_{it}E_{it}^{T} - c_{it}\right] \tag{3}$$

where A are beginning of period assets, R is the interest factor, w the hourly wage rate, h a fixed number of hours (corresponding to 500 hours per quarter), τ_w a proportional tax rate that is used to finance social insurance programs, F the fixed cost of work,[13] B_{it} unemployment benefits, T_{it} the monetary value of food stamps received, D_{it} the amount of disability insurance payments obtained, and E_{it}^{UI}, E_{it}^{DI}, and E_{it}^{T} are recipiency $\{0,1\}$ indicators for unemployment insurance, disability insurance, and the Food Stamps program, respectively.[14] Note also that there are costs to applying for disability insurance which we discuss below. We assume that individuals are unable to borrow:

$$A_{it} \geq 0$$

In practice, this constraint has bite because it precludes borrowing against unemployment insurance, against disability insurance, against social security and against the means-tested program.

At retirement, people collect social security benefits which are paid according to a formula similar to the one we observe in reality (see below). These benefits, along with assets that people have voluntarily accumulated over their working years, are used to finance consumption during retirement.

Unemployment Insurance We assume that unemployment benefits are paid only for the quarter immediately following job destruction. We define eligibility for unemployment insurance E_{it}^{UI} to mirror current legislation: benefits are paid only to people who have worked in the previous period, and only to those who had their job destroyed (job quitters are therefore ineligible for UI payments, and we assume this can be perfectly monitored).[15] We assume $B_{it} = b \times w_{it-1}\overline{h}$, subject to a cap, and we set the replacement ratio $b = 75\%$. This replacement ratio is set at this high

[13]The fixed cost of work is a pecuniary proxy for the disutility of work.

[14]We assume that food stamps are paid in cash rather than in the form of coupons. While this is in contrast with the reality, it would be of little practical importance if stamps were inframarginal or if there were "trafficking". Moffitt (1989) finds evidence for both phenomena.

[15]We have simplified considerably the actual eligibility rules observed in the US. A majority of states have eligibility rules which are tougher than the rule we impose, both in terms of the number of quarters necessary to be eligible for any UI and in terms of the number of quarters of work necessary to be eligible for the maximum duration (Meyer, 2002). However, making eligibility more stringent in our model is numerically difficult because the history of employment would become a state variable. Our assumption on eligibility shows UI in its most generous light.

value because the payment that is made is intended to be of a similar magnitude to the maximum available to someone becoming unemployed.

In the US, unemployment benefit provides insurance against job loss and insurance against not finding a new job. However, under current legislation benefits are only provided up to 26 weeks (corresponding to two periods of our model) and so insurance against not finding a new job is limited. Our assumption is that there is no insurance against the possibility of not receiving a job offer after job loss. This simplifying assumption means that, since the period of choice is one quarter, unemployment benefit is like a lump-sum payment to those who exogenously lose their job and so does not distort the choice about whether or not to accept a new job offer. The only distortion is introduced by the tax on wages.

Universal Means-Tested Program In modelling the universal means-tested program, our intention was to mirror partially the actual food stamps program but with three important differences. First, the means-testing is only on household income rather than on income and assets; second, the program provides a cash benefit rather than a benefit in kind; and third, we assume there is 100% take-up.[16] These assumptions mean the program plays the role of providing a floor to income for all individuals. This is similar to Hubbard, Skinner and Zeldes (1995). Gross income is given by

$$y_{it}^{gross} = w_{it} h P_{it} + \left(B_{it} E_{it}^{UI}\left(1 - E_{it}^{DI}\right) + D_{it} E_{it}^{DI}\right)\left(1 - P_{it}\right) \tag{4}$$

giving net income as $y = (1 - \tau_w) y^{gross} - d$, where d is the standard deduction that people are entitled to when computing net income for the purpose of determining food stamp allowances. The value of the program is then given by

$$T_{it} = \begin{cases} \overline{T} - 0.3 \times y_{it} & \text{if } E_{it}^{T} = 1 \ \left(\text{i.e., if } y_{it} \leq \underline{y}\right) \\ 0 & \text{otherwise} \end{cases} \tag{5}$$

The maximum value of the payment, $\overline{T}$, is set assuming a household with two adults and two children, although in our model there is only one earner. The term $\underline{y}$ should be interpreted as a poverty line. In the actual food stamp program, only people with net earnings below the poverty line are eligible for benefits ($E_{it}^{T} = 1$).

[16] The difficulty with allowing for an asset test in our model is that there is only one sort of asset which individuals use for retirement saving as well as for short-term smoothing. In reality, the asset test applies only to liquid wealth and thus excludes pension wealth (as well as real estate wealth and other durables).

Disability Benefits and Social Security A final element of the budget constraint is disability insurance. We assume that workers may find themselves in circumstances that would lead them to apply for disability insurance. First we allow only individuals who face a negative productivity shock to apply for disability. The requirement of a negative shock to wages is meant to mimic a health shock that induces permanent inability to work. Second, we require people to remain unemployed for at least one quarter before being able to apply for disability insurance and then they must remain unemployed in the quarter that the application is made. Again, this is meant to reflect the actual rules of the system (there is a waiting period of 5 months between application and receipt of benefits, and during this period the individual must be unemployed). Third, we assume that only workers above the age of 50 are eligible to apply for disability benefits.[17]

Conditional on applying for benefits, individuals have a fixed probability of obtaining the benefit which we obtain from actual data (50%, see Bound et al., 2004). If successful, the individual remains eligible for the rest of their working life and disability insurance becomes an absorbing state. If not successful, the individual has to remain unemployed another quarter before taking up a job. Individuals can only re-apply in a subsequent unemployment spell. The combination of disability and the means-tested program turns out to be very important in fitting the declining labor force participation profiles with age. Disability payments can provide a high replacement rate which is not affected by the duration of unemployment. However, the requirement that individuals spend two quarters unemployed before the disability application is resolved would discourage a large proportion of applicants were it not for the means-tested program which provides a floor to income during this application process.

The value of disability insurance is given by

$$D_{it} = \begin{cases} 0.9 \times \overline{w}_i & \text{if } \overline{w}_i \leq a_1 \\ 0.9 \times a_1 + 0.32 \times (\overline{w}_i - a_1) & \text{if } a_1 < \overline{w}_i \leq a_2 \\ 0.9 \times a_1 + 0.32 \times (a_2 - a_1) + 0.15 \times (\overline{w}_i - a_2) & \text{if } a_2 < \overline{w}_i \leq a_3 \\ 0.9 \times a_1 + 0.32 \times (a_2 - a_1) + 0.15\,(a_3 - a_2) & \text{if } \overline{w}_i > a_3 \end{cases} \tag{6}$$

where $\overline{w}_i$ is average earnings computed before the time of the application and a_1, a_2, and a_3 are thresholds we take from the legislation. We assume $\overline{w}_i$ can be approximated by the value of the permanent wage at the time of the application. Whether an individual is eligible (i.e., $E_{it}^{DI} = 1$)

[17] Interestingly, this was an actual requirement of the program at the time of inception (1956). In our model, it reflects the fact that health shocks triggering disability are rare before this age.

depends on the decision to apply ($DI_{it} = 1$) while being out of work and on having received a large negative productivity shock. We assume that the probability of success is independent of age. Eligibility does not depend on whether an individual quits or the job is destroyed.

By contrast with our assumption of a 50% probability of success for DI is our assumption of 100% take-up for our universal means-tested program and for unemployment insurance. We assume that this difference arises because of the difficulty of verifying disability compared to the income test and the unemployment test.

In retirement, all individuals receive social security calculated using the same formula used for disability insurance.

2.5 Employment Risk and Wage Risk

There are a number of sources of risk in our model. These include shocks to wages, job destruction, the rate at which job offers are sampled and the heterogeneity of firms. The shocks to wages are productivity risk. In a perfectly competitive labor market with no search frictions and no firm heterogeneity there is effectively no distinction between wage risk and employment risk. The unemployed are those who have received negative productivity shocks such that their productivity is below their reservation wage and so the individual prefers unemployment. In itself this does not constitute employment risk since it is an endogenous decision motivated by low earnings. Thus in the absence of labor market frictions, the distinction is meaningless.

The distinction between employment and productivity risk becomes relevant in the presence of search frictions and/or firm heterogeneity. Job destruction clearly leads to unwanted unemployment. However, if there were no uncertainty about receiving a new offer and no firm heterogeneity, there would be no employment risk as such: jobs would be located instantaneously and whether or not they are acceptable would depend only on individual productivity.

The presence of firm heterogeneity and match specific effects, means that some job may be available with a match value that would lead to a wage worth taking for an unemployed individual. Search frictions however, make it hard to find such a job and creates uncertainty in the length of unemployment. Moreover heterogeneity generates an option value to waiting in the unemployment state if the job arrival rate when on the job (and therefore the likelihood to be matched with a high-wage firm) is lower than the job arrival rate when unemployed. The uncertainty generated because of the combination of firm heterogeneity and search frictions we refer to as employment risk.

The productivity shocks that we observe (due to health shocks, demographic shocks, etc.) are assumed to be uninsurable uncertainty. We assume that there is no commitment from the side of the firm so Harris-Holmstrom type contracts are not implementable. Further, we assume there is no private insurance market against employment risk. This incomplete markets set-up is consistent with results from Attanasio and Davis (1996) and others.

An issue of central importance is the real degree of uncertainty. Measured wages vary; our structure strips out all variability that can be attributed to measurement error. Moreover we do not take the entire unexplained variance (after measurement error) as representing uncertainty, but only the innovation to wages from one period to the next. Of course this innovation may represent variability which is fully anticipated by the individual but totally unpredictable from the researcher's information. One strategy to resolve this issue of differing information sets is to use consumption data to "solve backwards" as in Blundell and Preston (1998) and Cuhna, Heckman and Navarro (2004). However, the Blundell and Preston methodology is problematic in the presence of a discrete participation choice.

Individuals move between firms and this leads to variation in earnings. We do not consider this as productivity risk per se; this is variability in earnings that is the result of a choice made by the individual in response to an offer. There is ex ante uncertainty about what type of firm will make an offer and whether an offer is received, but the ability to move between firms does not have a downside (i.e., bad offers can be turned down) and we label the risk associated with mobility as employment risk.

If such mobility was not possible, this might increase the amount of insurance firms would be willing to offer because of greater worker commitment. Part of the contribution of this paper is to separate out the variability in earnings due to productivity shocks, the variability due to employment risk and the variability due to endogenous choices of workers.

2.6 Numerical Solution

There is no analytical solution for our model. Instead, the model must be solved numerically, beginning with the terminal condition on assets, and iterating backwards, solving at each age for the value functions conditional on work status. The solution method is discussed in more detail in the appendix. Here we describe the main features of the algorithm used.

We start by constructing the value functions for the individual when employed and when out of work. When employed, the state variables are $\left\{A_{it}, u_{it}, a_{ij(t_0)}\right\}$, corresponding to current assets, individual productivity and the match effect. The match effect is indexed by t_0, which is the date the job began.[18] When unemployed and not on disability, the state variables are $\left\{A_{it}, u_{it}, DI_t^{Elig}\right\}$, corresponding to current assets, individual productivity and an indicator of whether the individual is eligible to apply for disability in that period. When unemployed and receiving disability, the state variables are $\{A_{it}, D_{it}\}$ where D_{it} is the amount of disability benefit received defined by equation (6).

We consider first the value function for an employed individual. An employed individual in the next period will have the choice of quitting into unemployment, moving to a new job or staying with the firm. However if the job is destroyed the individual will have to move to unemployment. Thus the value function for an individual i who is working in period t is

$$V_t^e\left(A_{it}, u_{it}, a_{ij(t_0)}\right) = \tag{7}$$
$$\max_c \left\{ \begin{array}{c} U\left(c_{it}, P_{it}=1\right) + \\ \beta\delta E_t\left[V_{t+1}^n\left(A_{it+1}, u_{it+1}, DI_{it+1}^{Elig}=1\right)\right] \\ +\beta\left(1-\delta\right)\left(1-\lambda^e\right)E_t\left[\max\left\{\begin{array}{c} V_{t+1}^n\left(A_{it+1}, u_{it+1}, DI_{it+1}^{Elig}=1\right), \\ V_{t+1}^e\left(A_{it+1}, u_{it+1}, a_{ij(t_0)}\right), \end{array}\right\}\right] \\ +\beta\left(1-\delta\right)\lambda^e E_t\left[\max\left\{\begin{array}{c} V_{t+1}^n\left(A_{it+1}, u_{it+1}, DI_{it+1}^{Elig}=1\right), \\ V_{t+1}^e\left(A_{it+1}, u_{it+1}, a_{ij(t_0)}\right), \\ V_{t+1}^e\left(A_{it+1}, u_{it+1}, a_{ij(t+1)}\right) \end{array}\right\}\right] \end{array} \right\}$$

The expectation operator is conditional on information at time t. If there is no offer available in $t+1$, the expectation operator is over the productivity shock only; if an offer is an offer in $t+1$, the expectation taken in t is also over the type of the firm making the offer.

We consider now the value function for an unemployed individual. Among the unemployed, we distinguish between those who have the option of applying for disability and those who are ineligible to apply (either because the individual is under 50 or because he has had an application turned

[18]Ideally we should model the behaviour of the firm. If the firm has a fixed number of positions, and if there are firing costs, a firm with characteristic $a_{ij(.)}$ may not make an offer to any worker. High $a_{ij(.)}$ firms may wish to wait to locate high u_{it} workers, in the same way that high u_{it} workers may wish to wait for high $a_{ij(.)}$ firms. At present we ignore this issue.

down in the previous period). The value function when ineligible for disability is given by:

$$V_t^n\left(A_{it}, u_{it}, DI_{it}^{Elig}=0\right) = \tag{8}$$
$$\max_c \left\{ \begin{array}{c} u\left(c_{it}, P_{it}=0\right) \\ \\ +\beta\lambda^n E_t\left[\max\left\{ \begin{array}{c} V_{t+1}^n\left(A_{it+1}, u_{it+1}, DI_{it}^{Elig}=0\right), \\ V_{t+1}^e\left(A_{it+1}, u_{it+1}, a_{ij(t+1)}\right) \end{array} \right\}\right] \\ \\ +\beta\left(1-\lambda^n\right) E_t\left[V_{t+1}^n\left(A_{it+1}, u_{it+1}, DI_{it}^{Elig}=0\right)\right] \end{array} \right\}$$

The expectation operator is again conditional on whether an offer has been received: if no offer has been received the only remaining uncertainty is over productivity.

For an individual who is eligible to apply for disability, the value function is given by

$$V_t^n\left(A_{it}, u_{it}, DI^{Elig}=1\right) = \max_{c,App}\left\{ u\left(c_{it}, P_{it}=0\right) + \beta \left\{ \begin{array}{ll} V_{t+1}^A & if\ Apply=1 \\ V_{t+1}^{NA} & if\ Apply=0 \end{array} \right. \right\} \tag{9}$$

where

$$\begin{aligned} V_{t+1}^{NA} &= \lambda^n E_t\left[\max\left\{ \begin{array}{c} V_{t+1}^n\left(A_{it+1}, u_{it+1}, DI^{Elig}=1\right), \\ V_{t+1}^e\left(A_{it+1}, u_{it+1}, a_{ij(t+1)}\right) \end{array} \right\}\right] \\ &\quad + \left(1-\lambda^n\right) E_t\left[V_{t+1}^n\left(A_{it+1}, u_{it+1}, DI^{Elig}=1\right)\right] \\ V_{t+1}^A &= S \times V_{t+1}^{DI}\left(A_{it+1}, D_{it+1}\right) + (1-S) \times E_t\left[V_{t+1}^n\left(A_{it+1}, u_{it+1}, DI^{Elig}=0\right)\right] \end{aligned}$$

and S is the exogenous probability of a successful application. When deciding whether or not to apply, the individual already knows if he has a job offer in that period. If the disability application is successful, we can calculate the resulting value function, V_{t+1}^{DI}, analytically: the amount of the disability insurance payment, D_{it}, depends on the permanent wage only and not on the particular firm that the individual has most recently been working for. This amount is earned each year until retirement.

In each period the individual decides, based on a comparison of these value functions, whether or not to work; and if working, whether or not to move to another job if the opportunity arises; and if not working, whether or not to apply for disability benefit. The decision about whether or not to move to another job if an outside offer is received is, in practice, more straightforward than the other decisions because we assume that there is no cost of switching firm. This means that the decision to switch firm involves a simple comparison of the $a_{ij(.)}$ and the individual will move if the

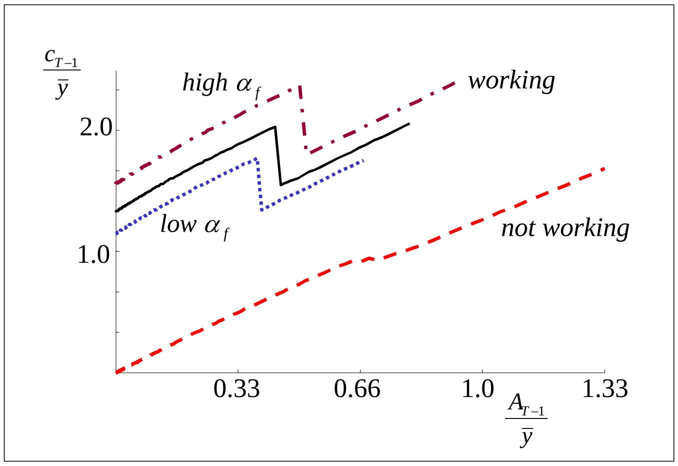

Figure 1: Consumption as a function of current assets conditional on current period work status

new offer is from a higher $a_{ij(.)}$-firm than the current one.[19]

The solution of the model consists of policy functions for consumption, participation, etc. Before turning to the results, it is instructive to illustrate part of the solution for the model. To give an example where discrete jumps in policy functions are more clearly illustrated, we show the solution without retirement and so the life-cycle ends at age 62. The same qualitative pictures are observed with retirement. Figure 1 shows consumption as a function of assets in the period preceding the end of life, $T-1$, for participants and non-participants, and for different firm types, conditioning on individual productivity. The point to stress here is that consumption is not monotonic in the asset stock even when conditioning on labor market status: this is because labor market status in future periods changes as the asset stock increases. For example, the sharp declines in consumption when participating at a given firm in $T-1$ arise at the asset stock which induces the individual not to work in period T. Because the individual is not working in period T, lifetime income is lower and consumption falls in both periods. On the other hand, since leisure is higher in the next period, overall welfare is higher: the value function is monotonically increasing in wealth. The extent of the fall depends on the degree of non-separability between consumption and leisure in the utility function: if consumption and leisure are closer substitutes, the discrete fall is less. However, the presence of the kinks arises because of the discreteness in the labor supply decision. If we look at

[19]If we were to allow for a cost of switching firm in the numerical solution, then the decision about whether or not to switch would depend on a comparison of the value function at the existing firm and the value function at the new firm. This difference will depend on the expected duration of the new job, the worker's horizon and all elements of the dynamic programing problem.

the solution in earlier time periods or the solution with retirement included, these sharp kinks are smoothed out. This is partly because the fall in income associated with a change in participation in one period in the future can be smoothed out over several periods. It is also partly because uncertainty smooths the discreteness: a marginal increase in asset holdings in period t will only change participation in $t+1$ in particular states and so has less of an impact on consumption in period t than if participation in $t+1$ changed in all states.

3 Data

We use the 1993 panel of the Survey of Income and Program Participation (SIPP) to estimate our wage dynamics parameters, and the 1988-1996 Panel Study of Income Dynamics (PSID) to construct participation and unemployment duration profiles. In both data sets, we stratify the sample by education, low (those with at least a high school diploma, but no college degree), and high (those with a college degree or more). The SIPP data have the advantage of giving information on wages around job switches. However, the short length of the SIPP panel means that it is not useful for duration analysis, so we use PSID data for that purpose.

3.1 The SIPP

The main objective of the Survey of Income and Program Participation (SIPP), conducted by the US Census Bureau, is to provide accurate and comprehensive information about the income and welfare program participation of individuals and households in the United States. The SIPP offers detailed information on cash and noncash income on a sub-annual basis. The survey also collects data on taxes, assets, liabilities, and participation in government transfer programs.

The SIPP is a nationally representative sample of individuals 15 years of age and older living in households in the civilian non-institutionalized population. Those individuals, along with others who subsequently come to live with them, are interviewed once every 4 months for a certain number of times (from a minimum of 3 to a maximum of 13 times, see below). Each year, a new "panel" starts, so some overlapping is expected. The first sample, the 1984 Panel, began interviews in October 1983 and surveyed individuals for 9 times. The second sample, the 1985 Panel, began in February 1985 and surveyed individuals for 8 times. We use the 1993 panel, which has 9 interviews in total (or three years of data for those completing all interviews).[20]

The Census Bureau randomly assigns people in each panel to four rotation groups. Each rotation group is interviewed in a separate month. Four rotation groups thus constitutes one cycle, called a wave, of interviewing for the entire panel. At each interview, respondents are asked to provide information covering the 4 months since the previous interview. The 4-month span is the reference period for the interview.

Our sample selection is as follows. The raw data has 62,721 records, one for each individual, corresponding to 1,767,748 month/person observations (note that, due to attrition, not all individuals complete 9 interviews). We drop females, those aged below 25 or above 60, those completing less than 9 interviews, the self-employed, those who are recalled by their previous employer after a separation, those with missing information about the state of residence, and some outliers in earnings.[21] Our final sample includes 6,226 individuals corresponding to 224,136 month-person observations, or 3 years of data per individual. We report some sample statistics in Table 6 in the appendix.

Our measure of (firm-specific) hourly wage is obtained by dividing annual earnings earned at the firm by annual hours worked at the firm. Individuals may have multiple hourly wage observations within a year if they work for multiple firms (concurrently or not). We use only the job that pays the highest proportion of annual earnings. In the SIPP, each job (firm) an individual is working for is assigned an ID. We set $M_{it} = 1$ if the employer the individual is working for at time t is different than the one he was working for at time $t-1$. We allocate individuals to the low and high education groups based on response to a question about the highest grade of school attended. An important advantage of the SIPP over the PSID when it comes to estimating the wage process allowing for job mobility is that the SIPP does not average pay over different employers. Thus the full effect of a move from one employer to another is observed.

3.2 The PSID

The PSID data are drawn from the 1988-1996 family and individual-merged files. The PSID started in 1968 collecting information on a sample of roughly 5,000 households. Of these, about 3,000 were representative of the US population as a whole (the core sample), and about 2,000 were low-income

[20] The raw data can be obtained at http://www.nber.org/data/sipp.html.

[21] An outlier is defined as one whose (annualized) earnings fall by more than 75% or grow by more than 250%.

families (the Census Bureau's SEO sample). Thereafter, both the original families and their split-offs (children of the original family forming a family of their own) have been followed. In the empirical analysis we use the core sample after 1988 because detailed data on monthly employment status and other variables of interest are available only after that year.

Our sample selection is as follows. We focus on males with no missing records on race, education, or state of residence. We drop those with topcoded wages, the self-employed, those with less than three years of data, and those with missing records on the monthly employment status question. Education level is computed using the PSID variable with the same name.

The PSID asked individuals to report their employment status in each month of the previous calendar year and their year of retirement (if any). We use these questions to construct a quarterly participation indicator for each individual and unemployment durations. We classify as not employed in a given month those who report to be unemployed/temporarily laid off, out of the labor force, or both, in that month. We treat unemployment and out-of-labor force as the same state; this tallies with the definition of unemployment that we use in the simulations (see Flinn and Heckman, 1991, for a discussion of the difference between these two reported states).[22] In principle, the durations are both left- and right-censored. Some spells begin before the time of the first interview, while some spells are still in progress at the time of the last interview. To avoid problems of left censoring we only use spells that begin in the sample. In calculating durations, we take our sample to be individuals who exit between 1988 and 1992. However, we use more recent years of PSID data (1993-1996) to calculate durations for those whose spells are right-censored by the 1988-1992 window. This reduces the censoring from 13.09% of all spells to 5.52%.

4 Estimating the Wage Process

Wages are observed conditional on individuals working; within-firm wage growth, which identifies the variance of permanent productivity shocks, is only observed if the individual does not change job; between firm wage growth, which helps identify heterogeneity across firms is observed only for job movers. Further, employment and mobility decisions are all endogenous and if this is ignored we risk biasing the estimates of the variances to wages and of firm heterogeneity.

[22] If the distinction in the data between out-of-labor force and unemployment reflects a difference in search intensity, we could make a meaningful distinction in our model only if we introduced a search decision with a cost attached.

To address this problem our approach is as follows: First we model the selection process into and out of work and between firms. We then construct sample selection terms and estimate wage growth equations conditioning on these terms. We finally obtain the estimates of the variances of interest by modelling the first and second moments of unexplained wage growth for various subgroups. We simplify the problem by assuming normality of all error terms.

Define the latent utility from labor market participation as $P_{it}^{*} = z_{it}'\varphi^{ed} + \pi_{it}$. The associated labor market participation index is $P_{it} = 1\{P_{it}^{*} > 0\}$, which is unity for participants. Workers separate from their current employer voluntarily (quits) or involuntarily (layoffs). As argued by Borjas and Rosen (1980), job turnover, regardless of who initiates it, represents the same underlying phenomenon, that of workers' marginal product being higher elsewhere. Let $M_{it}^{*} = k_{it}'\theta^{ed} + \mu_{it}$ denote the latent utility from moving in period t to an employer that is different from the one in period $t-1$ (this approximates the utility from moving to another firm). The indicator $M_{it} = 1\{M_{it}^{*} > 0\}$ singles out the "movers". We assume: $\begin{pmatrix} \pi_{it} & \pi_{it-1} & \mu_{it} \end{pmatrix}' \sim N(\mathbf{0}, \mathbf{I})$.

Taking first differences of the wage equation (1), using the process for permanent shocks (2) and recalling that $\xi_{it} = \left(a_{ij(t)} - a_{ij(t_0)}\right)$, we obtain:

$$\Delta \ln w_{it} = \Delta d_t^{ed} + \Delta x_{it}'\psi^{ed} + \zeta_{it} + \Delta e_{it} + \xi_{it} M_{it}$$

Wage growth is only observed for those who work in both periods. To achieve identification of the relevant parameters, we make the following assumptions (omitting for simplicity of notation the education-specific superscripts):

1. $E\left(a_{ij(t)}a_{ij(s)}\right) = \sigma_a^2$ if $j(s) = j(t)$ and zero otherwise.

2. We denote $\sigma_\zeta^2 = E\left(\zeta_{it}^2\right)$ and $\sigma_e^2 = E\left(e_{it}^2\right)$ (for all i, t) the variances of the permanent productivity shock and measurement error, respectively. We denote $E(\zeta_{it}\pi_{is}) = \sigma_\zeta \rho_{\zeta\pi}$ if $s = t$ and assume it to be zero otherwise.[23]

3. Given the definition of the mobility premium ξ_{it}, we assume $E(\xi_{it}\pi_{is}) = \sigma_\xi \rho_{\xi\pi}$ if $s = t$, $E(\xi_{it}\pi_{is}) = \sigma_\xi \rho_{\xi\pi_{-1}}$ if $s = t-1$, and zero otherwise.

[23] We denote with ρ_{ab} the correlation coefficient between a and b, and with σ_a the standard deviation of a.

4. We allow for contemporaneous correlation between the unobservable of the job mobility decisions (μ) and the shocks to the permanent productivity component and the match effect: $E\left(\zeta_{it}\mu_{is}\right) = \sigma_{\zeta}\rho_{\zeta\mu}$, and $E\left(\xi_{it}\mu_{is}\right) = \sigma_{\xi}\rho_{\xi\mu}$ for all $s = t$ and zero otherwise.

5. We assume that the distribution of innovations to the match effect ξ_{it} and the productivity shock are uncorrelated ($E(\xi_{it}\zeta_{is}) = 0\ \forall\ t, s$), and that there is no selection on measurement error ($E\left(e_{it}\pi_{is}\right) = E\left(e_{it}\mu_{is}\right) = 0\ \forall\ t, s$).

Suppose now that we select only those who work at t and $t-1$ ($P_{it} = 1, P_{it-1} = 1$). Using the law of iterated expectations is easy to show that:

$$\begin{aligned} E\left(\Delta \ln w_{it} | P_{it} = 1, P_{it-1} = 1\right) &= E\left(\Delta \ln w_{it} | P_{it} = 1, P_{it-1} = 1, M_{it} = 1\right) \Pr\left(M_{it} = 1\right) \\ &\quad + E\left(\Delta \ln w_{it} | P_{it} = 1, P_{it-1} = 1, M_{it} = 0\right)\left(1 - \Pr\left(M_{it} = 1\right)\right) \\ &= \Delta d_t^{ed} + \Delta x_{it}' \psi^{ed} + G_{it} \qquad (10) \end{aligned}$$

where G_{it} is a "selection" term induced by labor market participation in both periods and inter-firm mobility (see the Appendix for details).[24] The idea is to estimate the components of this selection term in a first stage (running separate probit regressions because of the assumed orthogonality assumption between π_{it} and μ_{it}), and use these to then estimate ψ^{ed} consistently in a second stage using only participants in both periods.

Define now unexplained wage growth (observed only for participants in both periods):

$$g_{it} = \Delta\left(\ln w_{it} - d_t^{ed} - x_{it}'\psi^{ed}\right) = \zeta_{it} + \Delta e_{it} + \xi_{it} M_{it} \qquad (11)$$

We can now use a method of moments procedure to identify the underlying stochastic process. The key parameters we need to identify are the variance of the permanent shocks and the variance of the firm level heterogeneity. We achieve this by using the first and second moments of the residuals for movers and for stayers, as well as the first-order autocovariance. In the process we not only estimate the two variances of interest but also all the relevant correlations that drive selection. The details of the moments we use are given in the Appendix.

[24] In estimation we do not use the restrictions on the parameters of interest imposed by (10). This only results in a loss of efficiency, but it does not affect consistency. We estimate the standard errors by the block bootstrap.

We adopt a multi-step estimation strategy. In a first step, we estimate probit regressions for labor market participation separately for each quarter, which is the assumed decision period. We also estimate a probit for mobility in period t conditioning on observing an individual working in both t and $t-1$. We set $M_{it}=1$ if this condition is satisfied and if the employer in period t differs from the one in period $t-1$.

In the second step we estimate (10) using only labor market participants in both periods. We define the hourly wage at firm j as annual earnings at firm j divided by annual hours of work at firm j. The selection term for participation for someone observed working a whole year and which is included in the construction of G_{it} is the aggregated quarterly Mills ratio, i.e. $\overline{\lambda_{\pi}} = \frac{1}{4}\sum_{q=1}^{4}\frac{\phi\left(z'_{it(q)}\varphi_q\right)}{\Phi\left(z'_{it(q)}\varphi_q\right)}$. This gives us estimates of ψ^{ed} and thus allows to construct consistent estimates of wage growth residuals g_{it}. In the final step, we estimate the structural parameters σ^2_{ζ}, σ^2_{ξ}, σ^2_{e}, and the various correlation coefficients. The variance of the match effect (σ^2_a) can be recovered from $\frac{\sigma^2_{\xi}}{2}=\sigma^2_a$. We consider a system of three non-linear equations for g_{it}, g^2_{it}, and $g_{it}g_{it-1}$, impose cross-equation constraints and estimate the three equations jointly by non-linear least squares.

Standard errors are computed using the block-bootstrap procedure suggested by Horowitz (2002). In this way we account for serial correlation of arbitrary form, heteroskedasticity, as well as for the fact that we use a multi-step estimation procedure, pre-estimated residuals and selection terms. We should point out that this procedure is likely conservative, since it allows for more serial correlation than that implied by the moment conditions we use. Thus p-values are likely upward biased.

4.1 Results

4.1.1 Participation and mobility

We start by estimating quarterly participation probits using the SIPP data. These include a quadratic in age, a dummy for whites, region dummies, a dummy for married, year dummies as well as unearned household income, and an index of generosity of the welfare system, which here we proxy with the generosity of the state-level UI system.[25] The latter two are excluded from the wage equation and are the instruments that identify selection into work - the unearned income as a pure

[25] To obtain a measure of the generosity of the UI program in the state where the worker lives, we rank states according to the ratio between maximum weekly UI benefit (which we take from current legislation) and average weekly wages (which we calculate from the CPS- using males only). Our measure of generosity is the rank variable, which varies over time and across states. We obtain similar results if we rank states pooling data for all years. Ideally, one would like to use an index of generosity of the Food Stamps program, but this is a federal program and its time-series varaibility is negligible.

income effect and UI as a fixed cost of work while eligible.[26] The participation equation for each quarter is reported in the appendix in Table 7. The main point is that unearned income has a strong and significantly negative effect on the probability of working. UI generosity is also a significant factor discouraging work, but only for the lower education group and not for the College graduates.

We also estimate a mobility probit, which will allow us to control for the censoring of the between firm wage growth. The dependent variable is whether an individual who was working in period t is in a different job in period $t+1$. Thus for the purposes of this estimation, mobility may include those moving jobs via unemployment.[27] The mobility probit includes the same variables as the participation equation, as well as industry dummies and an indicator as to whether the person was working for a non-profit organization, in both cases for period t. Unearned income influences positively mobility for both education groups; UI generosity influences positively mobility for the lower education group but not the College graduates. The effect of UI on mobility is theoretically ambiguous. On the one hand, it increases the reservation wage leading to individuals quitting employment following negative wage shocks and increasing mobility through this mechanism. On the other hand, when UI is low, durations of unemployment will be shorter and wage increases will occur through job-to-job mobility. Our results indicate that the former effect dominates. Our results also show that mobility declines with age for both groups. As people age, they tend to locate in better firms, and thus it becomes increasingly unlikely that an outside offer is sufficiently good to trigger mobility. Job destruction is an important force disrupting this age effect. The table with the results is also in the Appendix (Table 8).[28]

4.1.2 Variance Estimates

Armed with these results, we move on to estimate the parameters of the wage process by the method of moments, imposing constraints across equations. The moments we fit, together with the corrections for selection are reported in the appendix. The results are reported in Table 1.

[26] In practice we exploit variation over states and time. For the exclusion restrictions to be valid the US labour market should be sufficiently integrated and sufficient trade should be taking place, so that variability in benefits in one state does not affect the price of human capital in that state.

[27] We also tried to distinguish between "voluntary" movers (with no spell of unemployment in between two employment spells) and "involuntary" movers (those who move jobs via unemployment). We modified the selection process outlines in the Appendix, but find that the estimates of the variances of interest (σ^2_ζ and σ^2_α) change very little (although the correlation coefficient in Table 1 change considerably). Thus here we consider the simpler model with a single mobility index.

[28] Finally, unearned income and UI generosity as well as the industry dummies act as identifying instruments because they are excluded from the wage *growth* equation.

	Whole sample (1)	Low education (2)	High education (3)	Neglect selections (4)	Neglect mobility (5)	Neglect participation (6)
σ_ζ	0.114	0.104	0.116	0.147	0.149	0.111
	(0.016)	(0.032)	(0.015)	(0.013)	(0.013)	(0.016)
	[0%]	[0%]	[0%]	[0%]	[0%]	[0%]
σ_e	0.085	0.081	0.087	0.085	0.085	0.085
	(0.008)	(0.017)	(0.009)	(0.008)	(0.008)	(0.008)
	[0%]	[0%]	[0%]	[0%]	[0%]	[0%]
σ_a	0.213	0.208	0.215			0.209
	(0.013)	(0.026)	(0.019)			(0.012)
	[0%]	[0%]	[0%]			[0%]
$\rho_{\zeta\pi}$	0.217	0.193	-0.164		0.376	
	(0.184)	(0.193)	(0.269)		(0.146)	
	[8.8%]	[15.6%]	[61.6%]		[0%]	
$\rho_{\zeta\mu}$	-0.497	-0.901	-0.500			-0.713
	(0.330)	(0.506)	(0.409)			(0.292)
	[4.4%]	[1.2%]	[19.2%]			[0.8%]
$\rho_{\xi\pi}$	0.273	0.508	-0.416			
	(0.252)	(0.346)	(0.790)			
	[19.2%]	[3.6%]	[67.6%]			
$\rho_{\xi\pi_{-1}}$	-0.250	-0.253	0.098			
	(0.189)	(0.225)	(0.653)			
	[1.6%]	[10%]	[95.6%]			
$\rho_{\xi\mu}$	0.210	0.314	0.242			0.302
	(0.166)	(0.298)	(0.228)			(0.152)
	[11.6%]	[8%]	[20.8%]			[2%]

Note: σ_ζ, σ_e, and σ_a are the st.dev. of the permanent shock, measurement .error, and firm/matching effect. $\xi = a_j - a_{j-1}$. $\rho_{\zeta\pi}$ ($\rho_{\xi\pi}$) is the correlation between the permanent shock (mobility premium) and unobserved heterogeneity in the participation equation. $\rho_{\xi\pi_{-1}}$ is the correlation between the mobility premium and unobserved heterogeneity in the participation equation in the previous period. $\rho_{\zeta\mu}$ ($\rho_{\xi\mu}$) is the correlation between the permanent shock (mobility premium) and unobserved heterogeneity in the mobility equation. Standard errors (in parenthesis) are computed using the block bootstrap. Bootstrap p-values in square brackets.

Table 1: Wage variance estimates

The σ parameters refer to the standard deviations of the various stochastic components of wages. The ρ parameters are the correlations between the various stochastic shocks and the shocks driving selection. They are defined in appendix B. We estimate the model for the whole sample to have a comparison with previous work (column 1) and separately by the two education groups (columns 2 and 3).

Controlling for selection into employment and for job mobility, we find that in the whole sample the standard deviation of the permanent shock, σ_ζ, is about 0.11, the standard deviation of the transitory shock (measurement error), σ_e, 0.09, and the standard deviation of the firm shock, σ_a, 0.21. These parameters are all very precisely estimated. They imply a very important role of matching for wage dispersion; wages can fluctuate $\pm 42\%$ between firms for the same individual.

Columns (2) and (3) report the results of estimating the model separately for our two education groups. The stochastic process of wages is quite similar across the two education groups, although all variances are slightly higher for those with more education.

What happens if we ignore the fact that mobility is endogenous and attribute all wage fluctuations to the permanent and transitory shocks (σ_ζ and σ_e)? This, implicitly, has been the assumption made in papers estimating the covariance structure of earnings (MaCurdy, 1982; Abowd and Card, 1989; Meghir and Pistaferri, 2004) and in the precautionary savings papers estimating risk via the standard transitory/permanent shock decomposition (Carroll and Samwick, 2001; Gourinchas and Parker, 2002). In column (4) we report the results of this experiment as well as of not accounting for selection into work. We show that the estimated standard deviation of the permanent shock σ_ζ increases by about 30%. Whether this matters or not as far as the welfare implications of risk are concerned is an important focus of this paper and we consider this in our simulations.

To see the effect of ignoring selection in column (5) we ignore the mobility decision but account for the endogenous participation choice, while in column (6) we do the opposite. It is clear that what really matters is the firm mobility decision. Indeed, neglecting the participation correction reduces the variances of interest but the effects are minuscule.

We now turn to the estimated correlation coefficients between the random sources of wage growth and heterogeneity in the latent variables (ρ). In assessing these correlations it is important to note that the latent regressions are reduced forms, which do not include directly the wage; the heterogeneity will thus include components of the wage heterogeneity.

We assume that the permanent shock at time t is orthogonal to participation at time $t-1$, $cov\left(\zeta_{it}, \pi_{it-1}\right) = 0$, while $cov\left(\zeta_{it}, \pi_{it}\right) \neq 0$. People who receive a positive shock *should* generally be more likely to work. Consider an individual who has just had a boost to his productivity. If he quits into unemployment, he has no way of monetizing this, which would suggest $cov\left(\zeta_{it}, \pi_{it}\right) > 0$. However, since the shock is permanent, it may generate voluntary exit (this is similar to an income effect, although recall that there is a small probability that no new job offers will be received in the future). In the end, the effect is empirically ambiguous. In the whole sample, we estimate a positive effect, with bootstrap p-value 8.8%. However, this masks different effects between education groups, which appear not to be precisely estimated. When we reduce the number of estimated parameters by neglecting mobility (column 5) the correlation is positive and significant.

The correlation between the shock to mobility and the permanent shock to wages is harder to interpret mainly because in these equations, mobility includes job changes via unemployment. For job to job changes the shock to wages should be irrelevant because it is fully "portable". However, negative permanent shocks can induce entry into unemployment and hence movement to a new position. In this sense mobility and permanent shocks to wages can be negatively correlated. In fact $\rho_{\zeta\mu}$ turns out to be negative and mostly significant, except for the higher education group. In fact this is the group where there is less incidence of mobility via unemployment, so this result is consistent with the preceding argument.

The change in the "match effect" ξ may be correlated with the shocks to participation in both periods (π_{it}, π_{it-1}), and with the mobility shock μ_{it}. The correlation with μ_{it} is self-explanatory: mobility gains (high realizations of ξ) should be associated with a greater likelihood of observing a move, which would suggest $cov(\xi_{it}, \mu_{it}) > 0$. This is indeed what we find for the whole sample and for the two educational sub-samples.

The correlation between ξ_{it} and π_{it} (and π_{it-1}) is also relatively clear. Suppose one receives a high realization of ξ_{it}. If the individual quits, he loses the ability to step up on the ladder because good offers cannot be recalled (at least in the model). Thus one would expect $cov(\xi_{it}, \pi_{it}) > 0$ (and, symmetrically, $cov(\xi_{it}, \pi_{it-1}) < 0$). This is what we find in the whole sample and in the sub-sample of low educated individuals, while among the high educated the correlations have the opposite signs but are very imprecisely estimated.

These results show that a large amount of year-to-year wage variability is due to people moving to different firms and that ignoring this source of variability leads to wrong inferences regarding the extent of permanent productivity risk. This turns out to have important consequences for the welfare analysis of risk. This is partly because the various shocks underlying the types of risk have different persistence properties and individuals can react differently to their realizations.[29]

5 Calibrated Parameters

We now need to set the remaining parameters required for analysis. We impose values for some parameters such as the elasticity of intertemporal substitution and the discount rate using values

[29] While we have not imposed restrictions on the coefficients of the wage growth equation (10), we have checked whether the estimated coefficients are consistent with the structural estimates reported in Table 1, and found no violations.

from elsewhere in the literature. The rest we obtain through calibration using the structural model outlined in section 2.

We set the coefficient of relative risk aversion γ equal to 1.5, taken from Attanasio and Weber (1995), whose model of consumption also allows for nonseparable labor supply. The real interest rate is set equal to the real return on 3 month treasury bills, $r = 0.015$, and this is set equal to the discount rate $(\frac{1}{\beta} - 1)$.

In the model individuals have a 40 year working horizon (age 22-62) followed by a deterministic 10 year retirement spell. One period is assumed to be one quarter and so the model is solved for 160 periods when labor supply is chosen. A new job offer may be received each quarter, and similarly, the possibility of firm destruction is a quarterly event and decisions are taken each quarter. Further, each quarter individuals receive a productivity shock with probability 0.25 so productivity shocks occur on average once a year. This timing means individuals who stay with the same firm expect pay to be constant over a year. For the simulation we also require deterministic wage growth, because this defines the profile of life cycle income. To obtain this we use the workers in the PSID who do not move job, correcting for selection. Finally, the simulations take into account the social programs described earlier.

As described in the data section, the PSID asks individuals to report their employment status in each month of the previous calendar year. We use the answers to these questions to construct unemployment duration in quarters and a quarterly participation indicator for each individual. As said, we treat unemployment and out-of-labor force as the same state.

Given the estimated parameters of the wage process and those set above, we now set the remaining parameters to fit the life-cycle participation profile and unemployment duration profile for men, by education group. Our approach is to choose the parameters for each education group to minimise the sum of the absolute distance between statistics calculated in the data and corresponding simulated statistics. The statistics we use are the average participation rate and the median duration of unemployment in four age bands namely 22-31, 32-41, 42-51 and 52-61. In Table 9 in the appendix we show the fit of the moments we have targeted. In table 2, we present the calibrated parameter values. In Figures 2, 3, 4 and 5, we show the calibrated profiles.[30]

[30] Some of the differences between actual and fitted duration reflect time aggregation issues. In the model, people can be out of work only for an integer number of periods. In the data, we have durations in months and we use fractions when we convert this into quarters (i.e., a 4 month unemployment spell is equivalent to 1.33 quarter spell).

Parameter	High Education	Low Education
Job destruction rate δ	0.021	0.044
Job arrival rate - Unemployed λ^n	0.87	0.79
Job arrival rate - Employed λ^e	0.77	0.73
Fixed cost of work F	0.46	0.39
Disutility of participation η	-0.30	

Note: The values of δ, λ^n and λ^e are given as quarterly rates. We impose that the utility cost of participation η is the same across education groups.

The value of the fixed cost F
for each education group is expressed as a ratio to average earnings of that group at age 22.

Table 2: Parameters Obtained through Calibration

The job destruction rate is about twice as high for the lower educated individuals than for the higher educated ones. The contact rates are slightly higher for the more educated and they are higher for those out of work than when in work. The calibrated value of η is equivalent to consumption being 25.9% lower when the individual is participating. This value also implies consumption and leisure are substitutes, and thus it is consistent with the observed fall of consumption upon retirement (or unemployment). Finally, the fixed costs of work correspond to about 46% of average earnings for a 22 year old. For the low educated it is about 39%.

We now comment on the performance of the model both against the moments we match and against other properties of the model that we do not target. Figures 2 and 3 show participation profiles for the low educated and high educated. Each figure compares the profile in the data with the calibrated profile. For both education groups, participation rates are fairly constant until age 45, followed by a sharp decline to age 62. Part of this fall reflects early retirement, rather than temporary periods out of the labor force. Since early retirement is an endogenous labor supply response, we treat this in the same way as we treat unemployment. There is a level difference between the two groups: the high educated participate more than the low educated up to age 45 (participation rates around 96%, compared to 90% for the low educated), and the subsequent decline is less marked. Our match to participation is fairly good for both skill groups. This fitting of the down-turn in participation is achieved without allowing preferences or arrival rates to change with age. We also

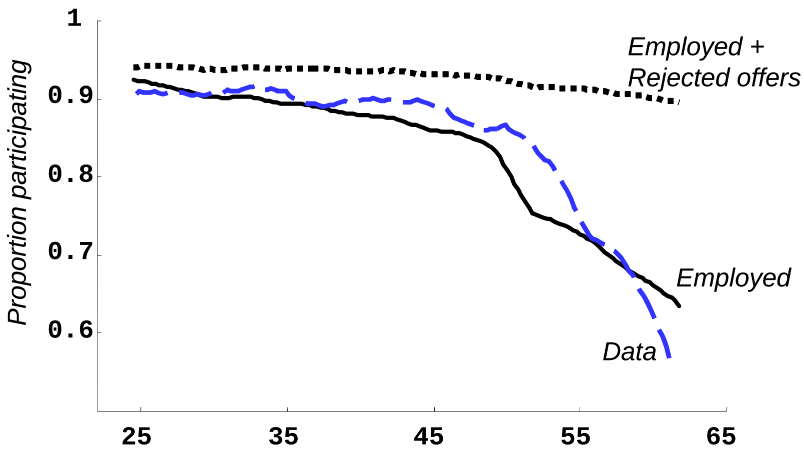

Figure 2: Actual and fitted participation profiles for the low education group

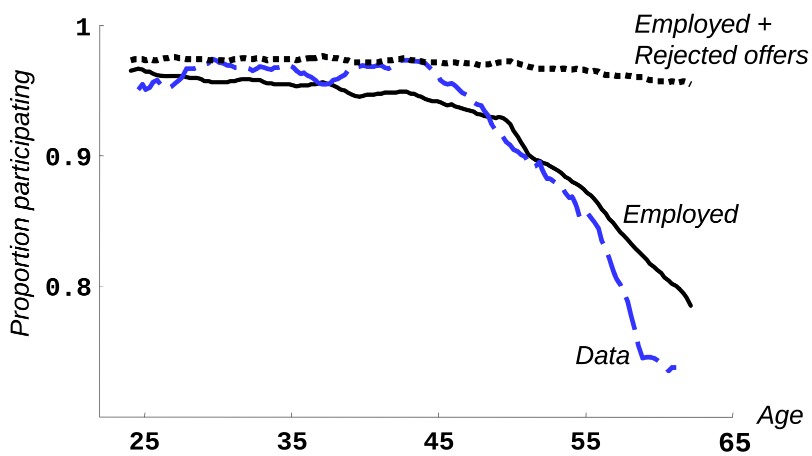

Figure 3: Actual and fitted participation profiles for the high education group

plot the participation rate that would be obtained if all offers received (including those offers from an existing employer) were accepted.

Figure 4 shows median duration over the life-cycle in the simulations and in the data for the high and low education groups. Durations have a maximum length determined by the number of quarters until age 62. In the data, durations are measured in months and are expressed as fractions of a quarter. In the simulations, durations are measured directly in quarters. Figure 5 shows comparable figures for mean durations. The duration data is skewed with the mean lying above the median particularly for older individuals. We do not attempt to match mean durations because of the problem of right-censoring in the data, particularly among the old. However, both mean and

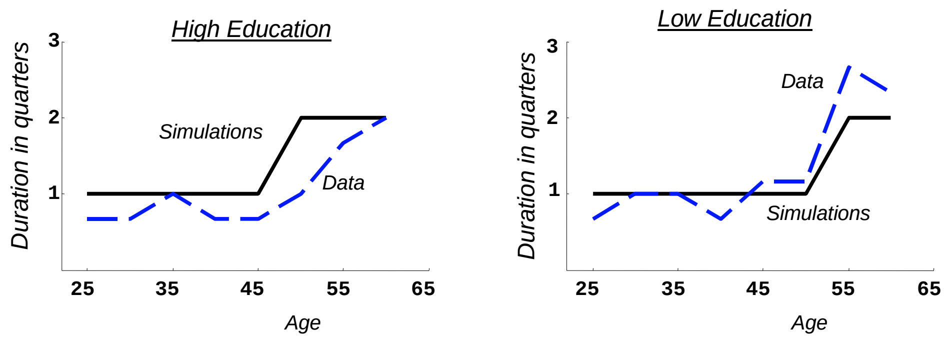

Figure 4: Actual and Fitted Median Durations by Education

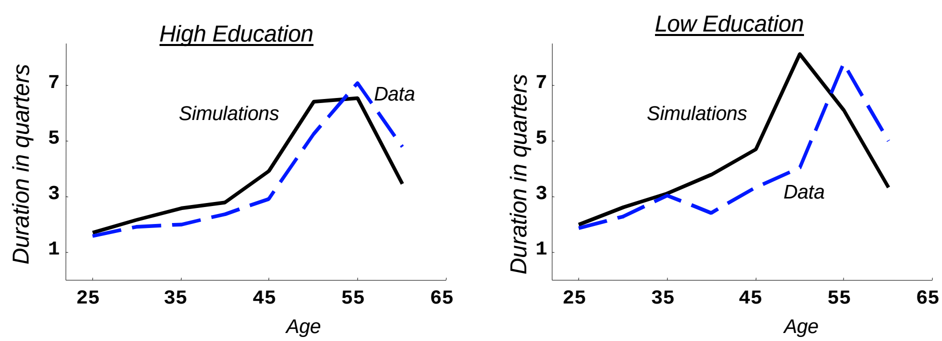

Figure 5: Actual and Fitted Mean Durations by Education

median simulated durations are fairly close to the corresponding paths in the data.

We have assumed $\lambda^n, \lambda^e, \delta, F$ and η are independent of age and so the age effects that we find in the simulated profiles can be explained only by endogenous saving and labor supply behavior in response to the budget constraint and the welfare benefit structure: the match in the slope of profiles over the life-cycle is not an artefact of age varying parameters and is a demonstration of the strength of the model.

5.1 Implications of the model

We have calibrated the model using only participation and unemployment duration data. However, the model has implications for a range of different variables. In particular, we use the model to predict the duration of employment ("tenure"), the wage loss associated with a spell of unemployment, the extent of consumption loss on unemployment and the arrival rate of accepted offers. Table 3 reports the model predictions and corresponding statistics in the data for a number of statistics that are

Statistic	Data	Model	
		High Education	Low Education
Median Duration of Employment			
Age $22 - 46$	12	27	15
Mean Consumption Loss (25-60)*			
$\Delta \ln c_{t+1}$	-0.14, -0.068	-0.172	-0.137
Mean Wage Loss+			
$\Delta \ln w_{t+1}$	-0.19	-0.227	-0.164
$\Delta \ln w_{t+4}$	-0.076	-0.108	-0.071
Arrival of accepted job offers:			
On-the-job (Age <40)		0.032	0.041
From unemployment (Age <40)		0.510	0.430

Employment durations are reported as number of quarters until employment with a particular employer ends. *The data numbers for consumption loss are taken from Browning and Crossley (2001) and Gruber (1997). +The data numbers for wage loss are taken from Jacobson, Lalonde and Sullivan (1993).

Table 3: Model Implications

not used in the calibration.

Employment durations The PSID has data on job tenure which could be used to pin down the arrival rate of offers while on the job, λ^e. Heads of household are asked how many years they have been working with their current employer. There are two main difficulties with these data. First, the spells are right-censored. Second, there is a substantial initial conditions problem with a very large proportion of spells being left censored. Third, a number of authors have questioned the reliability of these measures of reported tenure (Brown and Light, 1992). Because of these difficulties we do not use these data in the calibrations. Instead, we report statistics on employment duration from the data and from the simulations in Table 3. For the low educated, median employment duration in the simulations is 15 quarters for those younger than 46. For the high educated the median simulated spell length is 27 quarters. In the data, median employment duration is 12 quarters but this is likely to be downward biased because of censoring.

The Cost of Displacement There is empirical evidence that displaced workers experience earnings losses following job loss. Some authors impute this to exogenous skill depreciation during periods of unemployment (Rogerson and Schindler, 2001; Ljunqvist and Sargent, 2002). An alternative that is consistent with our model is that wages on re-entry may be lower than before job loss because of the loss of a particular good match on entering unemployment. We report in Table 3 the extent of the wage fall on re-entry.[31]

For the high educated, wages on re-entry are, on average, 22.7% lower than before displacement. For the low educated, the loss is 16.4%. These figures are similar to those found in the literature. In particular, we contrast these figures with those reported by Jacubson, LaLonde, and Sullivan (1993) for their non-mass layoff sample (after controlling for time trends). They report that 1 quarter after displacement, earnings of displaced workers are 19% less than before displacement. Finally, one implication of our model is that the displacement costs are likely to be short lived. Indeed, we calculate that 1 year after separation, wages of the low educated are only 7.1% of their pre-displacement wages. This figure is very close to the one we extrapolate from Jacubson, LaLonde, and Sullivan (1993) over a similar time horizon.

Consumption fall at unemployment Some recent papers have explored empirically the consumption loss associated with unemployment (Gruber, 1997; Browning and Crossley, 2001). In our model, the source of the loss will be the loss of the match as well as possibly the negative permanent shock. For comparison, we show simulated consumption loss in our model. In Table 3 we report average consumption loss by education group and compare to estimates in the literature. Consumption losses are higher for the high-educated (-17.2% vs. -13.7%) because means-tested government insurance programs are less effective for them.

Arrival rate of accepted offers on-the-job The final row of Table 3 reports the arrival rate of accepted offers among workers and among the unemployed. For workers, the arrival rate of accepted offers is low because workers only choose to move if they receive a better offer than the wage at their existing firm. The contrast between the high arrival rate of any offer and the low arrival rate of accepted offers is particularly striking. This is in stark contrast to the acceptance rate among the unemployed which is much higher in our simulations: there is less difference between the arrival rate of any offer (0.87 for the high educated) and the arrival rate of accepted offers (0.51).

[31] On a related matter, note that our model is consistent with the fact that low productivity may lead to unemployment. First, job destruction rates are higher for the low educated. Second, people who receive bad productivity shocks quit into unemployment.

6 Welfare Costs

One of the main aims of the paper is to show the extent to which different sources of risk matter for individual welfare. This is relevant particularly when evaluating policies such as unemployment insurance or the insurance value of tax credits, which effectively target part of the risk individuals face. In this model, we have exogenous, uninsured idiosyncratic shocks and so welfare will increase if insurance is provided. We also have behavioral responses to insurance built in both through changes in participation and through changes in savings. This means we can evaluate the risk sharing benefits of different sorts of insurance as well as identifying the behavioral effects induced by the insurance programs.

The model is partial equilibrium in that the wage process and interest rate are exogenous but we require the government budget to balance over the life cycle of a cohort, which is assumed to have N members. Thus we impose

$$\sum_{i=1}^{N}\sum_{t=1}^{T}\frac{1}{R^t}\left[\left(B_{it}E_{it}^{UI}\left(1-E_{it}^{DI}\right)+D_{it}E_{it}^{DI}\right)(1-P_{it})+E_{it}^{T}T_{it}\right]=\sum_{i=1}^{N}\sum_{t=1}^{T}\frac{1}{R^t}\tau_w w_{it}hP_{it}+Deficit \tag{12}$$

where B_{it} is unemployment insurance, D_{it} is disability insurance and T_{it} are food stamps; E_{it}^{UI}, E_{it}^{DI} and E_{it}^{T} are 1/0 indicators of eligibility for each of the programs respectively and $P_{it}=1$ denotes employment. On the right hand side $\tau_w w_{it}hP_{it}$ represents tax revenue from a working individual. The deficit term represents unaccounted expenditures and will be kept constant across all simulation experiments. Following a policy simulation we select the tax rate τ_w to satisfy this government budget constraint; individuals take τ_w as given.[32] Budget balance is imposed *within* a particular education group. We therefore abstract from the insurance *between* groups that Attanasio and Davis (1996) found to be important. Further, since there are no aggregate shocks in the economy and no business cycle fluctuations, we do not consider the value of, for example, smoothing the effect of the business cycle (Lucas, 1987; Storesletten et al., 2001). Thus we focus entirely on the cost to the individual of idiosyncratic risk, which would be smoothed out in a first best setting. Allowing the budget to balance over all education groups would confuse the issue we are considering with distributional questions.

[32] We assume that unemployment insurance and disability insurance are financed by the tax on wages, even though in reality the financing is partly imposed upon the firms. However, if the incidence of the tax falls on the workers, as most empirical studies find, our assumption is inconsequential.

To define the welfare cost of risk define the life time expected utility of an individual by

$$E_0 U_k = E_0 \sum_t \beta \frac{(c_{kt} \exp\{\eta P_{kt}\})^{1-\gamma}}{1-\gamma}$$

where the subscript k refers to the implied consumption stream in the baseline economy ($k = 1$) or an alternative economy with different risk characteristics ($k = 2$) and E_0 is the expectation at the beginning of working life. Now define π as the proportion of consumption an individual is willing to pay to face environment $k = 2$ rather than $k = 1$. This is implicitly defined by

$$E_0 U_2|_\pi \equiv E_0 \sum_t \beta \frac{(\pi c_{2t} \exp\{\eta P_{2t}\})^{1-\gamma}}{1-\gamma} = E_0 U_1$$

which implies that

$$\pi = \left[\frac{E_0 U_1}{E_0 U_2|_{\pi=1}}\right]^{\frac{1}{1-\gamma}}$$

We report values of π for small changes in risk. Specifically our welfare measure is defined by

$$\Delta_W = \frac{\pi}{\Delta\sigma/\sigma}$$

and is interpreted as the proportional change in consumption that an individual is willing to pay for a "small" change in the environment expressed as a proportional change ($\Delta\sigma/\sigma$) in some parameter σ (say a change in the standard deviation of a shock or in the rate of arrival of job offers).

Any change in the risk properties is likely to have labor supply effects, which can have an effect on equilibrium wages and on arrival rates. Here we have taken the route of considering small departures from the current environment. To consider impacts on a larger scale it would be important to model the firm side in some detail and allow for general equilibrium effects.

6.1 Welfare Cost of Risk

Table 4 shows results for the welfare elasticities with respect to the parameters determining risk. The first result that stands out is the effect of productivity risk (σ_ζ) on welfare relative to sources of employment risk: productivity risk is clearly the most serious source of uncertainty for individuals. High educated individuals are willing to pay 0.5% of consumption over the life cycle to obtain a 1% reduction in the standard deviation of the permanent shock to productivity. This compares to a willingness to pay only 0.07% to obtain a 1% reduction in job destruction. The low educated are willing to pay slightly less to avoid productivity risk, at 0.34%, but this is still substantially more than the 0.09% to avoid a 1% point reduction in the job destruction rate. The calculation of the cost of productivity risk differs from the usual calculations in the literature, such as by Storesletten et al. (2001), who calculate the welfare benefit of removing variation in the extent of idiosyncratic risk over the life-cycle. Such insurance removes heteroskedasticity but the risk to permanent productivity remains. Our calculations show that, compared to other sources of risk, it is this permanent risk to productivity which induces the greatest welfare loss. It is worth noting that part of the welfare loss arises because realizations of permanent shocks impact on retirement wealth with negative shocks reducing individuals' ability to save for retirement. The second point is that productivity risk has a greater welfare cost for the high education group: the difference in the *level* of income across the two education groups means that the universal means-tested program provides better insurance against bad productivity shocks for the low education group, and thus the low educated attach less cost to additional productivity risk relative to the higher educated.

The welfare numbers we report include the impact of the change in risk on total output, which can be substantial. We report, in the second column for each education group, the elasticity of output with respect to the change in risk. For example, with increases in job destruction lifetime wealth falls: for the high educated, a 1% increase in job destruction leads to a 0.12% fall in total output for the high education group and 0.18% for the low educated group. A fall in output will be caused partly by a reduction in time spent working and partly by individuals working at less productive times of their life (due to precautionary increases in labor supply) or in less productive matches. To the extent that the falls in output are caused by increased leisure this will increase utility, offsetting the fall caused by lower output.

Table 4 also shows the relative importance of different sources of employment risk, contrasting job destruction with arrival rate risk both on and off the job. Increases in the job arrival rate for the unemployed have a very small, almost negligible effect on welfare for both education groups. The effect on output is also negligible. This result is driven by the relatively small unemployment durations in the US which underpin our calibrated arrival rates. Job destruction has a larger effect because it disrupts matching to better jobs. Thus an individual who is displaced not only has to spend some time unemployed, but is unlikely to obtain as good a match as before. The importance of matching is reflected in the relative importance of the job arrival rate for the employed: an increased arrival rate is valued because it improves the rate at which individuals can obtain offers from better firms and indeed has a substantial effect on output. Because of the lower destruction rate of jobs among the higher education groups the welfare effect of increasing the rate of arrival of job offers for the employed is higher.

Finally it is interesting to consider the role of match heterogeneity. In models where labor supply and mobility are ignored this heterogeneity would translate into wage risk, because it would cause a change in wages associated with job moves. However, here such heterogeneity may have benefits for individuals since there is the chance of obtaining a better job offer while bad offers can always be turned down. Individuals in the high educated group are willing to pay 0.4% of consumption for a 1% increase in firm level heterogeneity. For the lower education group the figure is a little lower at 0.28%, partly because of the less stable employment profiles.

In the final row we show the welfare cost of productivity risk if there were no firm heterogeneity and all the variability due to job mobility were attributed to productivity risk. To achieve this we use results from estimating the deterministic growth in wages and the stochastic process for wage shocks ignoring job mobility. We then recalibrate the model under the assumption of no firm heterogeneity. The underlying numbers for the variance come from Column 5 in Table 1. The result is very interesting. For the high education group the elasticity of the willingness to pay for a 1% point reduction of σ_ζ increases for 0.49 to 0.6. The effect on output is also much larger, because the associated increase in the variance of wage shocks increases the number of periods out of work. For the lower education group, the welfare effects are less affected by this change. However, the output implications of increasing risk are now much higher. The means-tested program provides the welfare insurance, but this implies a direct increase in the number of periods out of work and an increase in taxes, which in turn discourage work. From these results its seems very important to consider carefully the sources of risk for welfare calculations and for the output effects of changing risk.

6.2 Welfare Benefit of Government Insurance

We now turn to an examination of the value of the various welfare programs we have included in our analysis. They provide insurance for different aspects of risk faced by the individual, although they are unlikely to provide anything close to full insurance.

We can use our model to assess the extent to which individuals value these programs in their

	High Education		Low Education	
Scenario	*Elasticity of consumption willing to pay (Δ_W)*	*Elasticity of Output*	*Elasticity of consumption willing to pay (Δ_W)*	*Elasticity of Output*
Productivity risk: σ_ζ	-0.490	-0.014	-0.340	-0.077
Job Destruction: δ	-0.074	-0.116	-0.087	-0.180
Unempl. arrival rate: λ^n	0.026	0.013	0.026	0.038
On-the-job arrival: λ^e	0.097	0.104	0.074	0.115
Firm heterogeneity: σ_a	0.396	0.425	0.279	0.365
Productivity risk without mobility	-0.607	-0.08	-0.394	-0.241

Table 4: Welfare effects of various sources of risk

current design. We follow the same approach of a local change as before. In this case, however, we consider a small (1%) change in the government deficit and compare the welfare effects of channelling this change into each of our programs in turn, namely UI, the universal means-tested program and DI. This calculation focuses on the insurance benefit of these programs because there is no cross-group redistribution. The results are presented in rows 1 to 3 of Table 5. Row 4 considers the welfare effect of channelling the increase in the deficit into a reduction in the proportional tax rate. For all the groups the most valuable program is the means-tested program because it provides some insurance against large negative (and permanent) shocks. High education individuals are willing to pay 0.25% of consumption for the increase in public expenditure to be channeled to the means-tested program. For the low educated this number is higher at 0.4%. The next most desirable program is UI. This is consistent with the earlier findings that seemed to suggest that productivity risk is more important than the risk of unemployment.

In considering the tax cut, the two groups are willing to pay 0.07% and 0.15% of consumption, respectively, to see the 1% increase in the deficit going to a tax cut. This implies both groups prefer the money to be spent on UI or the means-tested program, rather than on a decrease in taxation within their own group. This implication is no longer true when we consider the Disability Insurance program. While both value it somewhat, the higher education group is indifferent between spending the money on tax cuts or on DI and the lower education group would rather see tax cuts. This at

	High Education	Low Education
Scenario	*Elasticity of consumption willing to pay* (π)	*Elasticity of consumption willing to pay* (π)
Unemp. Insurance	0.170	0.227
Food stamps	0.253	0.399
Disability Insurance	0.070	0.127
Tax Change	-0.070	-0.151

Table 5: Welfare effects of government programmes

first sight surprising result can be explained by two factors. First there is substantial uncertainty as to whether one can obtain DI following a negative shock and there are important costs to applying; secondly there is a large moral hazard issue because a number of individuals who actually qualify would be able to return to work in future periods. Thus moral hazard and low insurance value combine to make this a relatively undesirable program.[33]

7 Conclusions

In this paper we have set up a model of employment and consumption over the life-cycle that allows for different sources of risk, which can be described as leading to productivity risk and employment risk. The model allows for job to job transitions as well as the more usual job to unemployment transitions. Within this context we estimate a wage process, which also allows for endogenous wage changes because of accepted job offers. We calibrate our model to obtain measures of job arrival and destruction rates that make the model consistent with observed participation rates and unemployment durations over the life-cycle. We use these measures to quantify the importance of productivity and employment risk.

We find that separately identifying the sources of risk and allowing for endogenous job-to-job movements changes substantially the welfare effects of risk. This arises firstly because the persistence properties of the different sources of risk are quite different: Productivity shocks are permanent,

[33] There are two caveats to these comments: first, these calculations ignore the interactions that may arise between increases in the tax rate needed to fund increased generosity of a program and the take-up of that program: the increased tax rate will make programs more valuable by reducing the benefit of being at work. Second, in practice these programs are funded by taxing the general population and consequently involve a large component of cross group insurance.

while shocks that lead to unemployment (such as job destruction) and to a change in match value are in effect transitory, with persistence depending on the job offer arrival rates for the employed and the unemployed. Secondly, a proportion of the observed variability in wages is due to choices individuals make about moving firm rather than due to uncertainty per se. In terms of the welfare cost of the components of risk, we find that the welfare benefits of reducing the job destruction rates or increasing job arrival rates are minimal. However productivity risk, as estimated from our wage process, has large welfare consequences.

A direct implication of our results is that safety net programs, such as Food Stamps, which are means-tested but offer longer duration insurance are highly valued relative to Unemployment Insurance. The low educated value unemployment insurance and food stamps more than the high education individuals, even if they have to fund it from taxation within their own group (rather than by transfers from wealthier individuals). By contrast, Disability Insurance, at least as modelled here, makes expansion of the program undesirable for both groups.

A Appendix: Numerical Solution

Households have a finite horizon and so the model is solved numerically by backward recursion from the terminal period. At each age we solve the value function and optimal policy rule, given the current state variables and the solution to the value function in the next period. This approach is standard. The complication in our model arises from the combination of a discrete choice (to participate or not) and a continuous choice (over saving). This combination means that the value function will not necessarily be concave. The discrete choice about whether to move or not is less problematic because we assume that there is no cost of moving. This means that the decision to move depends only on the relative size of the match effect in the current and new firm.

There are five state variables in this problem: age, employment status, the asset stock, the permanent component of earnings, u_{it}, and the match component, $a_{ij(t_0)}$. Age and employment status are both discrete. We also discretize both the permanent component of earnings and the distribution of possible matches, leaving the asset stock as the only continuous state variable. Since the permanent component of earnings is non-stationary, we are able to approximate this by a stationary, discrete process only because of the finite horizon of the process. We select the discrete nodes in this process to match the paths of the mean shock and the unconditional variance over the life-cycle. In particular, the unconditional variance of the permanent component must increase linearly with age, with the slope given by the conditional variance of the permanent shock.

Value functions are increasing in assets A_t but they are not necessarily concave, even if we condition on labor market status in t. The non-concavity arises because of changes in labor market status in future periods: the slope of the value function is given by the marginal utility of consumption, but this is not monotonic in the asset stock because consumption can decline as assets increase and expected labor market status in future periods changes. This problem is also discussed in Lentz and Tranaes (2001). By contrast, in Danforth (1979) employment is an absorbing state and so the conditional value function will be concave. Under certainty, the number of kinks in the conditional value function is given by the number of periods of life remaining. If there is enough uncertainty, then changes in work status in the future will be smoothed out leaving the expected value function concave: whether or not an individual will work in $t+1$ at a given A_t depends on the realization of shocks in $t+1$. Using uncertainty to avoid non-concavities is analogous to the use of lotteries elsewhere in the literature. In the value functions (7) and (**??**), the choice of participation status in $t+1$ is determined by the maximum of the conditional value functions in $t+1$.

In solving the maximization problem at a given point in the state space, we use a simple golden search method. We solve the model and do the calibration assuming this process is appropriate. We then check that the results in our baseline case are unaffected when we use a global optimizing routine, simulated annealing. It is worth stressing that there are parameter values for which the techniques we used do not work. In particular, as the variance of shocks gets sufficiently low, the non-concavities in the expected value functions become problematic.

B Appendix: Deriving Moments for the Variance of Wages

Wages are given by

$$\ln w_{it} = d_t + x_{it}'\psi + u_{it} + e_{it} + a_{ij(t_0)}$$

where $u_{it} = u_{it-1} + \zeta_{it}$ is the permanent component, e_{it} the measurement error, and $a_{ij(t_0)}$ is the match effect. For simplicity of notation, we omit the education-specific superscripts. Thus wage growth is

$$\Delta \ln w_{it} = \Delta d_t + \Delta x_{it}'\psi + \zeta_{it} + \Delta u_{it} + \xi_{it} M_{it}$$

where $\xi_{it} = \left(a_{ij(t)} - a_{ij(t_0)}\right)$. The latent indexes associated to working and moving are:

$$\begin{aligned} P_{it}^* &= z_{it}'\varphi + \pi_{it} \\ M_{it}^* &= k_{it}'\theta + \mu_{it} \end{aligned}$$

for all t. Note that conditioning on participation in periods t and $t-1$, and using the law of iterated expectations, we obtain:

$$\begin{aligned} E\left(\Delta \ln w_{it} | P_{it} = P_{it-1} = 1\right) &= E\left(\Delta \ln w_{it} | M_{it} = 0, P_{it} = P_{it-1} = 1\right)\left(1 - \Pr\left(M_{it} = 1\right)\right) \\ &\quad + E\left(\Delta \ln w_{it} | M_{it} = 1, P_{it} = P_{it-1} = 1\right)\Pr\left(M_{it} = 1\right) \\ &= \Delta d_t + \Delta x_{it}'\beta + G_{it} \end{aligned}$$

where

$$G_{it} = \rho_{\zeta\pi}\sigma_\zeta\lambda_{P=1} + \rho_{\xi\pi}\sigma_\xi\lambda_{P=1}\Phi\left(k_{it}'\theta\right) + \rho_{\xi\mu}\sigma_\xi\phi\left(k_{it}'\theta\right) + \rho_{\xi\pi_{-1}}\sigma_\xi\lambda_{P_{-1}=1}\Phi\left(k_{it}'\theta\right)$$

and $\lambda_{M=0} = \frac{\phi\left(k_{it}'\theta\right)}{1-\Phi\left(k_{it}'\theta\right)}, \lambda_{M=1} = \frac{\phi\left(k_{it}'\theta\right)}{\Phi\left(k_{it}'\theta\right)}$, $\lambda_{P=1} = \frac{\phi\left(z_{it}'\gamma\right)}{\Phi\left(z_{it}'\gamma\right)}$, and $\lambda_{P_{-1}=1} = \frac{\phi\left(z_{it-1}'\gamma\right)}{\Phi\left(z_{it-1}'\gamma\right)}$. Thus, G_{it} is a "selection term" accounting for conditioning on multiple indexes. Note that we do not exploit the restrictions on the coefficients on the selection terms. However, we check if they are satisfied once estimates of the structural parameters are obtained. The estimation of the equation above is standard (Heckman 2-step method).

The "structural" parameters (i.e., the variances of the wage shocks) are identified by the restrictions imposed on the moments of g_{it}. Using formulae from Tallis (1961), the first moment is:

$$\begin{aligned}
E\left(g_{it}|P_{it}=P_{it-1}=1,M_{it}=0\right) &= -\rho_{\zeta\mu}\sigma_{\zeta}\lambda_{M=0}+\rho_{\zeta\pi}\sigma_{\zeta}\lambda_{P=1} \\
E\left(g_{it}|P_{it}=P_{it-1}=1,M_{it}=1\right) &= \left(\rho_{\zeta\mu}\sigma_{\zeta}+\rho_{\xi\mu}\sigma_{\xi}\right)\lambda_{M=1}+\left(\rho_{\zeta\pi}\sigma_{\zeta}+\rho_{\xi\pi}\sigma_{\xi}\right)\lambda_{P=1}+\rho_{\xi\pi_{-1}}\sigma_{\xi}\lambda_{P_{-1}=1}
\end{aligned}$$

The parameters of the model are clearly not identified from the first moments alone. Consider then the second moment for workers that either stay or move:

$$E\left(g_{it}^2\,|P_{it}=P_{it-1}=1,M_{it}=0\right)=\sigma_{\zeta}^2\left(\begin{array}{c} 1-\rho_{\zeta\pi}^2 z_{it}'\gamma\lambda_{P=1}+\rho_{\zeta\mu}^2 k_{it}'\theta\lambda_{M=0} \\ -2\rho_{\zeta\pi}\rho_{\zeta\mu}\lambda_{P=1}\lambda_{M=0}\end{array}\right)+2\sigma_e^2$$

and

$$\begin{aligned}
E\left(g_{it}^2\,|P_{it}=P_{it-1}=1,M_{it}=1\right) &= \sigma_{\xi}^2\left(\begin{array}{c} 1-\frac{1}{2}\rho_{a\pi}^2 z_{it}'\gamma\lambda_{P=1}-\frac{1}{2}\rho_{a\pi}^2 z_{it-1}'\gamma\lambda_{P_{-1}=1}-2\rho_{a\mu}^2 k_{it}'\theta\lambda_{M=1} \\ +2\rho_{a\mu}\rho_{a\pi}\lambda_{M=1}\lambda_{P=1}-2\rho_{a\mu}\rho_{a\pi}\lambda_{M=1}\lambda_{P_{-1}=1} \\ -\rho_{a\pi}^2\lambda_{P=1}\lambda_{P_{-1}=1}\end{array}\right) \\
&\quad +\sigma_{\zeta}^2\left(\begin{array}{c} 1-\rho_{\zeta\pi}^2 z_{it}'\gamma\lambda_{P=1}-\rho_{\zeta\mu}^2 k_{it}'\theta\lambda_{M=1} \\ +2\rho_{\zeta\mu}\rho_{\zeta\pi}\lambda_{M=1}\lambda_{P=1}\end{array}\right)+2\sigma_e^2
\end{aligned}$$

Finally, we consider the first order autocovariance $E\left(g_{it}g_{it-1}\,|.\right)$. At least in principle, we could use information on those who work for three periods in a row and classify them on the basis of their mobility decisions. In practice, there are too few observations in the relevant categories to be able to get structural identification in this case. We thus assume $\Pr\left(M_t=1,M_{t-1}=1\right)\approx 0$ and consider only the restrictions on the unconditional autocovariance, namely

$$E\left(g_{it}g_{it-1}\right)=-\sigma_e^2$$

C Appendix: Tables on Participation, Mobility, Calibration

Variable	Mean	Standard deviation
Average hourly wage	14.75	7.33
Age	41.24	8.63
White	0.89	0.32
Married	0.78	0.42
Unearned income		
High education	0.59	0.49
Public sector	0.08	0.28
Northeast	0.20	0.40
North Central	0.28	0.45
South	0.25	0.43

Table 6: Summary Statistics, SIPP 1993 panel

	High school or less				College dropout or more			
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
Age	0.0205	0.0181	0.0184	0.0212	0.0109	0.0084	0.0084	0.0128
	(0.0042)	(0.0041)	(0.0042)	(0.0039)	(0.0020)	(0.0020)	(0.0020)	(0.0020)
$Age^2/100$	−0.0307	−0.0282	−0.0286	−0.0317	−0.0151	−0.0121	−0.0122	−0.0176
	(0.0049)	(0.0048)	(0.0048)	(0.0045)	(0.0024)	(0.0023)	(0.0023)	(0.0023)
White	0.1012	0.0987	0.0913	0.0880	0.0414	0.0494	0.0455	0.0422
	(0.0134)	(0.0133)	(0.0131)	(0.0122)	(0.0086)	(0.0088)	(0.0086)	(0.0082)
Married	0.1426	0.1568	0.1570	0.1537	0.0539	0.0512	0.0512	0.0610
	(0.0105)	(0.0105)	(0.0105)	(0.0097)	(0.0063)	(0.0062)	(0.0063)	(0.0062)
Region dummies	21.11	30.02	29.01	34.68	12.11	12.01	6.89	17.85
	(3, 0%)	(3, 0%)	(3, 0%)	(3, 0%)	(3, 0.7%)	(3, 0.7%)	(3, 8%)	(3, 0%)
Year dummies	0.65	0.69	0.07	2.21	0.59	0.04	0.49	6.07
	(2; 72%)	(2; 76%)	(2; 97%)	(2; 33%)	(2; 74%)	(2; 98%)	(2; 78%)	(2; 5%)
Unearned income	−0.0442	−0.0411	−0.0376	−0.0442	−0.0110	−0.0107	−0.0116	−0.0114
	(0.0019)	(0.0018)	(0.0017)	(0.0019)	(0.0007)	(0.0006)	(0.0007)	(0.0007)
UI generosity	−0.0009	−0.0009	−0.0009	−0.0006	0.0002	−0.0000	−0.0000	−0.0000
	(0.0004)	(0.0004)	(0.0004)	(0.0004)	(0.0002)	(0.0002)	(0.0002)	(0.0002)

Note: The table reports marginal effects. Asymptotic standard errors in parenthesis. For region and year dummies we report the value of the χ^2 statistics of joint significance and, in parenthesis, the degrees of freedom and the p-value of the test.

Table 7: The Participation Decision

	High school or less	College dropout or more
Age	−0.0144 (0.0035)	−0.0109 (0.0029)
$Age^2/100$	0.0144 (0.0041)	0.0096 (0.0035)
White	−0.0150 (0.0109)	0.0095 (0.0079)
Married	0.0081 (0.0076)	−0.0103 (0.0066)
Not-for-profit	−0.0451 (0.0147)	0.0234 (0.0147)
Industry dummies	85.44 (4 df; p-value χ^2 0%)	66.62 (4 df; p-value χ^2 0%)
Region dummies	3.68 (3 df; p-value χ^2 30%)	5.66 (3 df; p-value χ^2 13%)
Year dummies	33.74 (2 df; p-value χ^2 0%)	77.58 (2 df; p-value χ^2 0%)
Unearned income	0.0025 (0.0005)	0.0012 (0.0003)
UI generosity	0.0008 (0.0003)	−0.0002 (0.0003)

Note: The table reports marginal effects. Asymptotic standard errors in parenthesis. For region and year dummies we report the value of the χ^2 statistics of joint significance and, in parenthesis, the degrees of freedom and the p-value of the test.

Table 8: The Mobility Decision

Statistic	High Education		Low Education	
	Data	Model	Data	Model
Participation Rate Age $22-31$	0.96	0.96	0.91	0.91
Participation Rate Age $32-41$	0.97	0.95	0.90	0.89
Participation Rate Age $42-51$	0.93	0.93	0.87	0.83
Participation Rate Age $52-61$	0.79	0.84	0.68	0.69
Median Duration Age $22-31$	1.00	1.00	0.67	1.00
Median Duration Age $32-41$	0.67	1.00	1.00	1.00
Median Duration Age $42-51$	1.17	1.00	1.00	1.00
Median Duration Age $52-61$	2.50	2.00	1.83	2.00

Table 9: Observed and Matched Moments

References

[1] Abowd, J. and D. Card (1989), “On the covariance structure of earnings and hours changes”, Econometrica 57, 411-45.

[2] Abowd, J.M., Kramarz, F. and D.N. Margolis (1999) “High Wage Workers and High Wage Firms” Econometrica, 67(2), 251-333.

[3] Attanasio, O., J. Banks, C. Meghir, and G. Weber (1999), “Humps and Bumps in Lifetime Consumption”, Journal of Business and Economic Statistics, 17(1), 22-35.

[4] Attanasio, O. and Davis, S. (1996) “Relative wage movements and the distribution of consumption” Journal of Political Economy 104: 1227-1262

[5] Attanasio, O., and G. Weber (1995), “Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey”, Journal of Political Economy, 103(6), 1121-57.

[6] Banks, J., R. Blundell, and A. Brugiavini (2001), “Risk pooling, precautionary saving and consumption growth”, Review of Economic Studies.

[7] Blundell, R., and I. Preston (1998), “Consumption inequality and income uncertainty”, Quarterly Journal of Economics 113, 603-640.

[8] Borjas, G.J. and S. Rosen (1980) “Income prospects and job mobility of younger men,” in R. Ehrenberg (ed.), Research in Labor Markets, Vol.3. Greenwich, CT: JAI Press, 159-18I.

[9] Bound, J., J. Cullen, A. Nichols, L. G. Schmidt (2004), “The Welfare Implications of Increasing Disability Insurance Benefit Generosity”, Journal of Public Economics, 88 (12): 2487-2514.

[10] Brown, J.N. and A. Light (1992), “Interpreting Panel Data on Job Tenure”, Journal of Labor Economics, 10(3), 219-257.

[11] Browning, M. and T. Crossley (2001) ”Unemployment Benefit Levels and Consumption Changes” Journal of Public Economics, 80(1), 1-23

[12] Caballero, R. (1990), “Earnings Uncertainty and Aggregate Wealth Accumulation”, The American Economic Review, 81(4), 859-871.

[13] Carroll, C. (1992), "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence", Brookings Papers on Economic Activity 1992(2), 61-156.

[14] Carroll C., and A. Samwick (1997), "The Nature of Precautionary Wealth", Journal of Monetary Economics 40.

[15] Costain, J. (1997), "Unemployment Insurance with Endogenous Search Intensity and Precautionary Saving", mimeo.

[16] Cuhna, F., J. J. Heckman and S. Navarro (2005) Separating Uncertainty from Heterogeneity in Life Cycle Earnings," NBER Working Papers 11024

[17] Danforth, J. P. (1979) "On the Role of Consumption and Decreasing Absolute Risk Aversion in the Theory of Job Search" in: Studies in the Economics of Search, ed. by S. A. Lippman, and J. J. McCall, pp. 109-131. North-Holland, New York.

[18] Deaton, A. (1991) "Saving and liquidity constraints" Econometrica 59(5), 1121-1142

[19] Farber H. and R. Gibbons (1996) "Learning and Wage Dynamics." Quarterly Journal of Economics, 111(4): 1007-1047

[20] Flinn, C. and J. Heckman (1983), "Are Unemployment and Out of the Labor Force Behaviorally Distinct Labor Force States?", Journal of Labor Economics, 1(1), 28-42.

[21] French, E., (2005) "The effects of health, wealth and wages on labour supply and retirement behavior" Review of Economic Studies 72(2): 395-427

[22] Gourinchas, P., and J. Parker (2002), "Consumption over the Life-Cycle", Econometrica 70(1): 47-89

[23] Gruber, J. (1997), "The Consumption Smoothing Benefits of Unemployment Insurance", The American Economic Review, Vol. 87(1),192-205.

[24] Horowitz, J. (2002) "Bootstrap Methods for Markov Processes," Econometrica 71: 1049-1082.

[25] Hubbard, R.G., J. Skinner, and S. Zeldes (1995), "Precautionary Saving and Social Insurance," Journal of Political Economy, 103(2), pages 360-99.

[26] Jacubson, L. S., LaLonde, R. J. and Sullivan, D.G. (1993) "Earnings Losses of Displaced Workers." The American Economic Review, 83(4), pp. 685-709.

[27] Kimball, M. (1990), "Precautionary Savings in the Small and in the Large", Econometrica 58, 53-73

[28] Lentz, R. (2003) "Optimal Unemployment Insurance in an Estimated Job Search Model with Savings" Boston University, mimeo

[29] Lentz, R. and T. Tranaes (2005), "Job Search and Savings: Wealth Effects and Duration Dependence", Journal of Labor Economics, 23(3), 467-490.

[30] Ljungqvist, L. and T. Sargent (1998), "The European Unemployment Dilemma", The Journal of Political Economy, 106(3), 514-550.

[31] Low, H. (2005) "Self-insurance and unemployment benefit in a life-cycle model of labor supply and savings", Review of Economic Dynamics 8(4): 945-975.

[32] Lucas, R.E. (1987) Models of Business Cycles Basil Blackwell, New York

[33] MaCurdy, T. (1982) "The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis" Journal of Econometrics 18(1): 83-114

[34] Manski, C. (2004), "Measuring Expectations," Econometrica, Econometric Society, vol. 72(5), pages 1329-1376.

[35] Meghir, C. and L. Pistaferri (2004) "Income variance dynamics and heterogeneity " Econometrica.

[36] Meyer, B. (2002) "Unemployment and workers' compensation programmes: rationale, design, labour supply and income support" Fiscal Studies 23: 1-49

[37] Moffitt, R. (1989), "Estimating the value of an in-kind transfer: The case of Food Stamps", Econometrica 57: 385-409

[38] Moffitt, R. and P. Gottschalk (1995), "Trends in the Covariance Structure of Earnings in the U.S.: 1969-1987", unpublished.

[39] Pistaferri, L. (2001), "Superior information, income shocks, and the permanent income hypothesis", Review of Economics and Statistics.

[40] Pistaferri, L. (2003), "Anticipated and unanticipated wage shocks, wage risk, and intertemporal labor supply", Journal of Labor Economics.

[41] Postel-Vinay, F., J.-M. Robin (2002) "Equilibrium Wage Dispersion with Heterogeneous Workers and Firms" Econometrica, 70: 2295-1350

[42] Rendon, S. (2006), "Job Search and Asset Accumulation under Borrowing Constraints", International Economic Review, 47(1), 233-263.

[43] Rogerson, R. and M. Schindler (2002) "The welfare costs of worker displacemen" Journal of Monetary Economics, vol. 49(6), pages 1213-1234

[44] Storesletten, K., Telmer, C., and A. Yaron (2001) "The welfare cost of business cycles revisited: Finite lives and cyclical variation in idiosyncratic risk" European Economic Reivew 45, 1311-1339

[45] Tallis G.M. (1961), "The Moment Generating Function of the Truncated Multi-normal Distribution", Journal of the Royal Statistical Society. Series B (Methodological), 23(1), 223-229.

[46] Topel, R.H. (1991) "Specific Capital, Mobility, and Wages: Wages Rise with Job Seniority" Journal of Political Economy, 99(1): 145-176

[47] Topel, R.H, and M.P. Ward (1992) "Job Mobility and the Careers of Young Men" Quarterly Journal of Economics, 107(2): 439-79.

[48] Zeldes, S. P., (1989). "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," The Quarterly Journal of Economics, vol. 104(2), pages 275-98.